

A Statistical Approach to Set Classification by Feature Selection with Applications to Classification of Histopathology Images

Sungkyu Jung*

Department of Statistics

University of Pittsburgh

Pittsburgh, Pennsylvania 15260, U.S.A.

E-mail: sungkyu@pitt.edu

Xingye Qiao

Department of Mathematical Sciences

Binghamton University, State University of New York

Binghamton, New York 13902-6000, U.S.A.

E-mail: qiao@math.binghamton.edu

*Corresponding author

Abstract

Set classification problems arise when classification tasks are based on sets of observations as opposed to individual observations. In set classification, a classification rule is trained with N sets of observations, where each set is labeled with class information, and the prediction of a class label is performed also with a set of observations. Data sets for set classification appear, for example, in diagnostics of disease based on multiple cell nucleus images from a single tissue. Relevant statistical models for set classification are introduced, which motivate a set classification framework based on context-free feature extraction. By understanding a set of observations as an empirical distribution, we employ a data-driven method to choose those features which contain information on location and major variation. In particular, the method of principal component analysis is used to extract the features of major variation. Multidimensional scaling is used to represent features as vector-valued points on which conventional classifiers can be applied. The proposed set classification approaches achieve better classification results than competing methods in a number of simulated data examples. The benefits of our method are demonstrated in an analysis of histopathology images of cell nuclei related to liver cancer.

KEYWORDS: Bioinformatics; Canonical angles; Discriminant analysis; Hotelling's T -square; Principal component analysis; Multidimensional scaling; Set classification.

1 Introduction

As advances in technology ease semi-automated segmentation and preprocessing of cell nucleus images, more pathologists are relying on histopathology to discriminate diseased tissues from benign tissues. The classification of tissues based on microscopic examination is often achieved using many cell nucleus images. Figure 1 illustrates two sets of cell nucleus images from human liver tissues. The eight nuclei in the left panel belong to a set labeled as normal tissues, while another set of eight nuclei from hepatoblastoma tissues is shown on the right. See Section 6 for the background and analysis of the data. An eminent statistical task to aid pathologists is to develop a method to classify a new tissue sample consisting of many nucleus images into the normal or the malignant. The classification rule in need is learned from sets of observations and also should be able to predict a single class label for a new set of observations. Such a problem, which we call *set classification*, has not been studied much in the statistical literature, although it appears to be useful in image-based pathology (Samsudin and Bradley, 2010, Wang et al., 2010).




Figure 1: Eight images of cell nuclei in a set labeled as normal tissue (a), and those from a set labeled as hepatoblastoma tissue (b). Each gray-scale image has 192×192 pixels. The set classification makes use of the set membership, to discriminate normal sets from hepatoblastoma sets.

To precisely define the set classification problem, suppose there are N tissue samples, each of which is represented by a set \mathcal{X}_i consisting of n_i images of cell nuclei \mathbf{x}_{ij} in the tissue, *i.e.*, $\mathcal{X}_i = \{\mathbf{x}_{i1}, \dots, \mathbf{x}_{in_i}\}$. Each set \mathcal{X}_i or the corresponding tissue sample has its label y_i (say, $y_i \in \{\text{normal}, \text{cancerous}\}$). Based on these N sets of cell nuclei images, we wish to

predict the label y^\dagger for a new tissue sample containing n_\dagger images, $\mathcal{X}_\dagger = \{\mathbf{x}_1^\dagger, \dots, \mathbf{x}_{n_\dagger}^\dagger\}$.

A few characteristics of such a data set make the task challenging. First, the order of observations $\{\mathbf{x}_{i1}, \dots, \mathbf{x}_{in_i}\}$ in each set \mathcal{X}_i may be only given by convenience, meaning that there is no correspondence between observations in different sets. In such a case, \mathbf{x}_{i1} and $\mathbf{x}_{\iota 1}$ in different sets ($i \neq \iota$) should not be directly compared. Moreover, the number of observations n_i in each set may be different from one another. These characteristics rule out a naive approach for set classification, which is based on a long vector consisting of the n_i observations $\mathbf{X}_i = [\mathbf{x}_{i1}^T, \dots, \mathbf{x}_{in_i}^T]^T$. Instead, a more appropriate statistical model is obtained when a set is regarded as an empirical distribution of observations.

To facilitate the understanding of the set as a distribution, a scatterplot of the liver cell nuclei data is overlaid with contours of the estimated normal densities in Figure 2. A point in Figure 2 corresponds to an image of cell nucleus and each is from a tissue labeled as normal or hepatoblastoma if marked by \times or \circ , respectively. Each observation in a set is assumed to be drawn from a distribution, whose mean and covariance are represented by the contour of the density. A useful insight of the data is that different sets have different distributional parameters. A visual inspection of the plot leads us to believe that the mean and covariance parameters will be useful for classification of sets. Classification based on the parameters of the distribution, or the features of the set, is the initial idea which we develop further in this paper.

Previous approaches to the classification of sets of images are mostly focused on extracting context-based features. For example, Tsantis et al. (2009) used morphological features and Wang et al. (2010) used shape and texture features. To the best knowledge of the authors, Ning and Karypis (2008) and Wang et al. (2010) were among the first to use set information in classification. Although not stated explicitly, they seem to have assumed a simple model of set classification, in that the set information of observations is not used in training but only in prediction.

Having examined more general set classification models, we are proposing a context-free




Figure 2: Scatters of cell nucleus images with their mean and covariance as ellipses for each of 10 sets, projected onto the first two principal component directions. This figure appears in color in the electronic version of this article.

feature selection approach to set classification. Our method is based on extracting statistical features of sets such as the mean and principal components (PCs; Jolliffe, 2002). The proposed feature extraction–selection method transforms empirical PC directions and subspaces into feature vectors so that conventional classifiers can be applied in the transformed space. Multidimensional scaling (Borg and Groenen, 2005) is extensively used in mapping PC subspaces into a real vector space. Our procedure is not a single method but a general framework which can be coupled with any off-the-shelf classifier. We demonstrate the use of variants of linear and quadratic discriminant analysis (Friedman, 1989, Srivastava and Kubokawa, 2007), Support Vector Machine (Cristianini and Shawe-Taylor, 2000) and Distance-Weighted Discrimination (Marron et al., 2007, Qiao et al., 2010) in this article. Prediction of the class label for a new observation is also based on its extracted features using multidimensional scaling.

Since we use only statistical features which are free of context, the proposed approach can be applied to many other types of data beyond images.

In the next section, we formulate relevant statistical models for set classification. We then propose our set classification framework in Section 3 and feature selection procedure in Section 4. A simulation study is presented in Section 5 to examine the effectiveness of our methods in high-dimensional situations. In Section 6, the proposed methods are illustrated

with a histopathological cell nucleus image data set. We conclude the article with a discussion in Section 7.

2 Statistical models for set classification

In this section we describe statistical models from which the set classification problem arises. In set classification, each observation \mathbf{x} is labeled by two indices: $y \in \{1, \dots, K\}$ for the class label and $i \in \{1, \dots, N\}$ for the set membership. The distribution of \mathbf{x} is different for different sets and classes. A natural model incorporating this characteristic of the data is a model with hierarchy, *i.e.*, a top level for the class label and a lower level for the set membership.

In the hierarchical model, a random observation \mathbf{x} follows a probability distribution $F(\mathbf{x} \mid \theta)$ where θ is the idiosyncratic parameter of each set. The parameter θ of a set is considered as a random variable whose distribution is determined by a different class label, y . That is, the dependence of \mathbf{x} on the class label y is defined only through its set membership. Denote Θ_i as a random parameter of the distribution of the i th set. The hierarchical model for set classification is

$$\begin{aligned} \Theta_i \mid (Y_i = k) &\sim H_k \quad (k = 1, \dots, K), \\ \mathcal{X}_i \mid (\Theta_i = \theta) &\sim f(\mathcal{X}_i; \theta) = \prod_{j=1}^{n_i} f(\mathbf{x}_{ij}; \theta) \quad (i = 1, \dots, N), \end{aligned} \tag{1}$$

where H_k is the distribution of the parameter θ for the k th class. In the model above, we have assumed that the observations $\mathbf{x}_{i1}, \dots, \mathbf{x}_{in_i}$ in a set \mathcal{X}_i are independent and follow an absolutely continuous distribution with density function f .

To elucidate the use of the hierarchical model, consider a model from which the data in Figure 2 may follow. Assuming normal distribution for each observation \mathbf{x}_{ij} , it is visually evident from Figure 2 that the sets (or the corresponding distributions) within either the normal or the hepatoblastoma group have different means and covariance matrices. The

difference, however, is small compared to the difference between the normal and hepatoblastoma groups. The (random) parameters $\Theta_i = (\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)$ of the i th set may be modeled as $\boldsymbol{\mu}_i | (Y = k) \sim N_p(\delta_k, \sigma^2 \mathbf{I}_p)$ and $\boldsymbol{\Sigma}_i | (Y = k) \sim m^{-1} \mathcal{W}_p(\Sigma_{(k)}, m)$, where $\mathcal{W}_p(A, m)$ denotes the Wishart distribution of dimension p with mean A and m degrees of freedom. The hyperparameters $(\delta_k, \Sigma_{(k)})$ differ by the class label $k \in \{\text{normal, hepatoblastoma}\}$. The set \mathcal{X}_i is then a set of i.i.d. observations following $N_p(\mu_i, \Sigma_i)$.

Other examples of the hierarchical model include situations where only the mean parameter $\boldsymbol{\mu}_i$ is random while the covariance parameter $\boldsymbol{\Sigma}_i$ is fixed, or vice versa. These models are exemplified in the top two panels of Figure 3. Their population structures, displayed as ellipses, clearly visualize the underlying hierarchical model. In both examples, classification of sets can be done by utilizing the parameters (features) of the distribution, represented as the locations and orientations of the ellipses. The separating hyperplane in Figure 3 by the linear discriminant analysis, pooling all observations in different sets together, is clearly less useful than using the distributional features.

A special case of the hierarchical model can be obtained by fixing the random Θ_i conditioned on $Y = k$ as a deterministic θ_k . This model is then no longer hierarchical and is referred to as a simple model. As illustrated in the bottom panels of Figure 3, all observations with the same class label have a common distribution. Since set membership plays no role here, a conventional classification method with the weighted voting strategy is Bayes optimal (see Web Appendix A). In our experience with simulated data analysis, our method based on the hierarchical model is superior to the weighted voting method, which is shown in Section 5.

3 Methods

Our framework for set classification is two-step. First, features of a set, or the parameters of the corresponding probability distribution, are estimated, selected and transformed. Then,

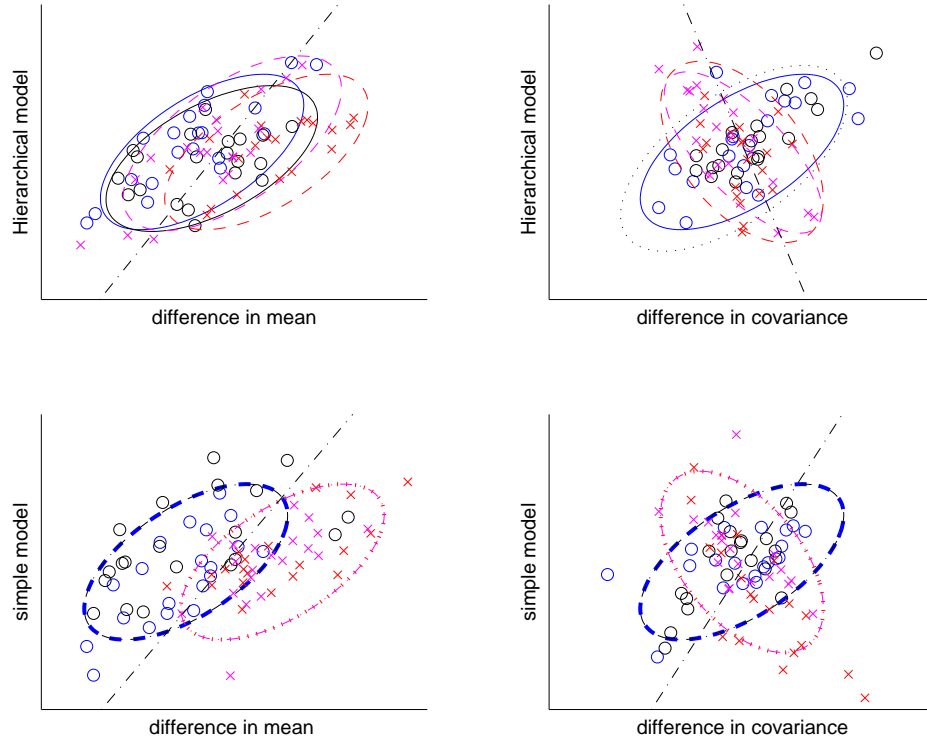


Figure 3: Toy data examples from the hierarchical (top) and the simple (bottom) models. Hierarchical models allow different distributions (depicted as the shapes of ellipse) for different sets, while the simple model ignores the set membership. A linear separating hyperplane may not be useful as can be seen in the right panels. This figure appears in color in the electronic version of this article.

a classifier is built based on the features extracted from the training data set.

3.1 Feature extraction

In choosing appropriate features for sets, it is worth revisiting the characteristics of the liver cell nuclei data illustrated in Figure 1 and 2. First, each observation in the data lies in a fairly high-dimensional space while the number of sets and the number of observations are relatively small. Modeling with a large number of parameters is problematic due to collinearity and can lead to overfitting. Second, as can be seen in Figure 2, the location and orientational information of the distribution (the principal axis of the ellipses) is useful to visually separate sets between different groups. These findings are what motivated us to consider using principal component analysis (PCA; Jolliffe, 2002) to learn orientational information as well as the sample mean for the location to extract succinct but useful features of a set.

In the following, we discuss a method of extracting the orientation of the distribution through PCA. An application of multidimensional scaling (MDS) follows in order to represent the principal component features in a low-dimensional vector space. The vector-valued features which contain orientational information are then attached to the sample means for use in the sequel.

3.1.1 Principal component spaces

Denote the sample mean of the observations in the i th set ($i = 1, \dots, N$) by $\hat{\boldsymbol{\mu}}_i = n_i^{-1} \sum_{j=1}^{n_i} \mathbf{x}_{ij} \in \mathbb{R}^p$. The empirical covariance matrix $\hat{\Sigma}_i \equiv n_i^{-1} \sum_{j=1}^{n_i} (\mathbf{x}_{ij} - \hat{\boldsymbol{\mu}}_i)(\mathbf{x}_{ij} - \hat{\boldsymbol{\mu}}_i)^T$ for the i th set is eigen-decomposed into $\hat{\Sigma}_i = \sum_{\ell=1}^p \lambda_{\ell i} \mathbf{e}_{\ell i} \mathbf{e}_{\ell i}^T$, where the eigenvectors (PC directions) $\{\mathbf{e}_{1i}, \dots, \mathbf{e}_{pi}\}$ are orthonormal to each other and the eigenvalues (PC variances) $\lambda_{1i} \geq \dots \geq \lambda_{pi} \geq 0$ are in descending order. In the high-dimensional, low-sample size situation, *i.e.*, $p \gg n_i$, we have $\lambda_{\ell i} = 0$ for all $\ell \geq n_i$.

The direction of the major variation is represented by the first few empirical PC directions

$\{\mathbf{e}_{1i}, \dots, \mathbf{e}_{ri}\}$ for some r . However, using the direction vectors as features can lead to an erroneous conclusion. Since the PC directions are axial, *e.g.*, $\mathbf{e}_{1i} = \mathbf{u}$ is essentially the same as $\mathbf{e}_{1i} = -\mathbf{u}$, a special care is needed. Moreover, denote $(\mathbf{e}_\ell^o, \lambda_\ell^o)$ for the eigen-pair from the population covariance matrix. The smaller $|\lambda_1^o - \lambda_2^o|$ is, the larger the variation in the empirical PC directions \mathbf{e}_{1i} and \mathbf{e}_{2i} . In the worst case, where $\lambda_1^o = \lambda_2^o$, the population PC directions \mathbf{e}_1^o and \mathbf{e}_2^o are indistinguishable and thus unidentifiable. In such a case, the estimates \mathbf{e}_{1i} and \mathbf{e}_{2i} are only meaningful if they are understood as a whole. To be more precise,

PROPOSITION 1. *If $\lambda_1^o = \lambda_2^o > \lambda_3^o$ then the subspace spanned by the first two population PC directions $\mathcal{L}^o = \text{span}(\mathbf{e}_1^o, \mathbf{e}_2^o)$ is identifiable, and $\mathcal{L}_i^{(2)} = \text{span}(\mathbf{e}_{1i}, \mathbf{e}_{2i})$ is a consistent estimator of \mathcal{L}^o .*

A proof of Proposition 1 can be found in Web Appendix A, in which the consistency of subspaces is formally defined.

On this end, denote the subspace spanned by the first r empirical PC directions of the i th set by $\mathcal{L}_i^{(r)} = \text{span}(\mathbf{e}_{1i}, \dots, \mathbf{e}_{ri})$, $1 \leq i \leq N$, for $1 \leq r \leq \min(p, n_i - 1)$. Which r to use is an important matter for model selection, and in Section 4 we propose a data-driven method to choose r . The arguments in the current section are applied to any given r .

Note that the collection $\{\mathcal{L}_i^{(r)} : i = 1, \dots, N\}$ does not lie in a Euclidean vector space, but in the Grassmannian manifold $G(r, p)$ (Chikuse, 2003). Therefore instead of analyzing elements in the curved manifold, $\mathcal{L}_i^{(r)}$ s are mapped onto a linear vector space. Such a mapping cannot be obtained by conventional dimension reduction methods such as PCA but can be achieved by multidimensional scaling. We first point out a distance function between $\mathcal{L}_i^{(r)}$ s, on which the mapping is defined.

3.1.2 Distances between principal component spaces

The distance between two subspaces of \mathbb{R}^p can be measured in terms of the canonical angles formed by the subspaces (*cf.* Stewart and Sun, 1990). As an intuitive example for the

canonical angles, consider a simple case where both subspaces are of dimension 1. The two subspaces are each spanned by unit vectors \mathbf{x}_1 and \mathbf{y}_1 . Then the distance between the subspaces is measured as the *smallest angle* θ formed by the basis vectors, *i.e.*, $\cos \theta = |\mathbf{x}_1^T \mathbf{y}_1|$, which is also called the canonical angle. For more general situations, canonical angles are computed as follows. Without loss of generality, let \mathcal{L}_1 and \mathcal{L}_2 be two subspaces with dimensions s and $r(\leq s)$ respectively, and \mathbf{L}_1 and \mathbf{L}_2 be matrices consisting of basis vectors of \mathcal{L}_1 and \mathcal{L}_2 , respectively. In particular, the $p \times r$ orthonormal matrix $\mathbf{L}_i = [\mathbf{e}_{1i}, \dots, \mathbf{e}_{ri}]$ is a basis matrix for the PC subspace $\mathcal{L}_i^{(r)}$. Let γ_ℓ be the ℓ th largest singular value of $\mathbf{L}_1^T \mathbf{L}_2$. Then the canonical angles between \mathcal{L}_1 and \mathcal{L}_2 are $\theta_\ell(\mathcal{L}_1, \mathcal{L}_2) = \arccos(\gamma_\ell)$ for $\ell = 1, \dots, r$.

A distance between \mathcal{L}_1 and \mathcal{L}_2 is now defined using the canonical angles. We use a modified Euclidean sine metric (*cf.* [Stewart and Sun, 1990](#)), defined by $\rho_s(\mathcal{L}_1, \mathcal{L}_2) = c(\sum_{\ell=1}^r \sin^2(\theta_\ell))^{\frac{1}{2}}$, where $c > 0$ is a constant. Since $\theta_\ell \in (0, \pi/2]$, $\rho_s(\mathcal{L}_1, \mathcal{L}_2)$ is at most $c\sqrt{r}$. The role of the constant c is to make the distance $\rho_s(\mathcal{L}_1, \mathcal{L}_2)$ commensurate with other features such as the mean. Specifically, when measuring pairwise distances among the PC spaces $\mathcal{L}_i^{(r)}$ s, we choose c to be the average of the empirical total variance in $\mathcal{L}_i^{(r)}$ s, $c = \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^r \lambda_{i\ell}$. This choice of c leads to good classification results, which are reported in Web Appendix B.

3.1.3 Mapping $\mathcal{L}_i^{(r)}$ via multidimensional scaling

Multidimensional scaling (MDS) maps N objects with distances $\rho_s(\mathcal{L}_i^{(r)}, \mathcal{L}_j^{(r)})$ into corresponding points $\mathbf{z}_1, \dots, \mathbf{z}_N$ in an m -dimensional space \mathbb{R}^m . The vector-valued configuration $\{\mathbf{z}_1, \dots, \mathbf{z}_N\}$ is found by minimizing a loss function, which measures how well the original distances $\rho_s(\mathcal{L}_i^{(r)}, \mathcal{L}_j^{(r)})$ are approximated by $\|\mathbf{z}_i - \mathbf{z}_j\|_2$. MDS can be used to map data in a non-Euclidean space into a vector space so that conventional multivariate methods (based on the Euclidean geometry) can be used. In extracting features corresponding to the PC spaces $\mathcal{L}_i^{(r)}$, we use MDS to construct a mapping from the non-Euclidean objects $\mathcal{L}_i^{(r)} \in G(r, p)$ into vector-valued points. We use the classical multidimensional scaling (CMDS) since it gives

an analytic solution requiring no numerical optimization. Detailed discussion on methods of MDS can be found in [Borg and Groenen \(2005\)](#) and [Cox and Cox \(2008\)](#). Note that PCA is not applicable for $\mathcal{L}_i^{(r)}$, which lies on a curved manifold $G(r, p)$.

Mapping training sample: In CMDS, the matrix of configuration $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_N]$ is obtained from the $N \times N$ matrix of pairwise squared-distances $\Delta = (\rho_s^2(\mathcal{L}_i^{(r)}, \mathcal{L}_j^{(r)}))_{i,j}$. Writing the matrix of pairwise squared Euclidean distances between columns of \mathbf{Z} as $D(\mathbf{Z}) = (\|\mathbf{z}_i - \mathbf{z}_j\|_2^2)_{i,j}$, we minimize the loss function

$$L(\mathbf{Z}) = \left\| -\frac{1}{2} C (D(\mathbf{Z}) - \Delta) C \right\|_F^2, \quad (2)$$

where $C = \mathbf{I}_N - \frac{1}{N} \mathbf{1}_N \mathbf{1}_N^T$ is the centering matrix, and $\|A\|_F$ is the Frobenius norm of matrix A . An analytic solution that minimizes (2) is available, due to [Gower \(1966\)](#).

THEOREM 2. *Let $B = -\frac{1}{2} C \Delta C$. From the eigen-decomposition of $B = Q \Lambda Q^T$, let Λ_+ be the diagonal matrix consisting of the positive eigenvalues and Q_+ the corresponding columns of Q . Then $\mathbf{Z} = \Lambda_+^{\frac{1}{2}} Q_+^T \in \mathbb{R}^{m \times N}$ minimizes (2).*

The configuration \mathbf{Z} lies in m -dimensional real vector space, where $m < \min(p, N)$ is the number of positive eigenvalues of B . Note that $B_{\mathbf{Z}} = \mathbf{Z}^T \mathbf{Z}$ is the $N \times N$ Gram matrix (consisting of the inner products of columns of \mathbf{Z}), and that if all eigenvalues of B are nonnegative, $B_{\mathbf{Z}}$ coincides with B . The coordinate matrix \mathbf{Z} is the mapped image in \mathbb{R}^m of the PC spaces $\mathcal{L}_i^{(r)}$, based on which a classifier can be trained.

Mapping for prediction: For a new observation \mathcal{X}_{\dagger} , we wish to extend the loss function (2) in order to map the new PC space $\mathcal{L}_{\dagger}^{(r)}$ for classification. We include the new observation in the loss function while the training points \mathbf{Z} are fixed. Define the $(N+1) \times (N+1)$ symmetric matrix of pairwise squared-distances

$$\Delta_{\dagger} = \begin{pmatrix} \Delta & \boldsymbol{\delta}_{12} \\ \boldsymbol{\delta}_{21} & 0 \end{pmatrix},$$

where $\boldsymbol{\delta}_{12} = (\boldsymbol{\delta}_{21})^T = (\rho_s^2(\mathcal{L}_i^{(r)}, \mathcal{L}_{\dagger}^{(r)}))_{i=1, \dots, N}$. Let $\mathbf{Z}_{\dagger}(\mathbf{z})$ be the $(m+1) \times (N+1)$ matrix of

coordinates, where the first N columns of $\mathbf{Z}_\dagger(\mathbf{z})$ represent the training points and the last column \mathbf{z} represents the new point. That is,

$$\mathbf{Z}_\dagger(\mathbf{z}) = \begin{pmatrix} \mathbf{Z} & \mathbf{z} \\ \mathbf{0}_N^T & \end{pmatrix},$$

where \mathbf{Z} is the minimizer of (2). While the training configuration \mathbf{Z} spans m dimensions, the new point \mathbf{z} may not be found in $\text{span}(\mathbf{Z})$. Therefore the dimension of \mathbf{Z}_\dagger is increased by 1. We generalize (2) to minimize

$$L_N(\mathbf{z}) = \left\| -\frac{1}{2}P_N(D(\mathbf{Z}_\dagger(\mathbf{z})) - \Delta_\dagger)P_N^T \right\|_F^2, \quad (3)$$

where $P_N = \mathbf{I}_{N+1} - \frac{1}{N}\mathbf{1}_{N+1}[\mathbf{1}_N^T; 0]^T$ is a projection matrix that works similarly to the centering matrix C except that the mean of the first N columns (*i.e.*, the training points \mathbf{Z}) is used for the centering.

For the purpose of classification, the last coordinate of the new $\mathbf{z} \in \mathbb{R}^{m+1}$ is not needed, because a classification rule learned from \mathbf{Z} only uses the first m coordinates. An analytic solution for the first m coordinates of \mathbf{z} , hereafter denoted by \mathbf{z}_\dagger , is also available. Define

$$B^\dagger = -\frac{1}{2}P_N\Delta_\dagger P_N^T = \begin{pmatrix} B & \mathbf{b}_{12} \\ \mathbf{b}_{21} & b_2 \end{pmatrix}.$$

It can be seen that \mathbf{b}_{12} is the vector of inner products $(\mathbf{z}_i^T \mathbf{z}_\dagger)_{i=1,\dots,N}$, *i.e.*, $\mathbf{b}_{12} = \mathbf{Z}^T \mathbf{z}_\dagger = Q_+ \Lambda_+^{\frac{1}{2}} \mathbf{z}_\dagger$. We then have the coordinates of the new point

$$\mathbf{z}_\dagger = \Lambda_+^{-\frac{1}{2}} Q_+^T \mathbf{b}_{12}. \quad (4)$$

The following theorem justifies the use of (4) as the feature vector of the new observation \mathcal{X}_\dagger .

THEOREM 3. *If $\mathbf{z}_\dagger^T \mathbf{z}_\dagger \leq b_2$, then*

$$\hat{\mathbf{z}} = \begin{pmatrix} \mathbf{z}_\dagger \\ \pm \sqrt{b_2 - \mathbf{z}_\dagger^T \mathbf{z}_\dagger} \end{pmatrix} \quad (5)$$

are the local minima of $L_N(\mathbf{z})$ (3). If the equality holds, i.e., $\mathbf{z}_\dagger^T \mathbf{z}_\dagger = b_2$, then (5) is the global minimum of $L_N(\mathbf{z})$. Moreover, $\mathbf{z}_\dagger^T \mathbf{z}_\dagger \leq b_2$ holds if B^\dagger is nonnegative definite.

To the best of our knowledge, Theorem 3 is the first attempt to obtain a closed form solution for mapping new observations onto coordinates of training configuration by CMDs. The loss function $L_N(\mathbf{z})$ is a fourth order polynomial of multiple arguments, thus finding a global minimum of L_N is a challenging task, involving iterative numerical methods (cf. Nie, 2006). Theorem 3 provides sufficient conditions for the existence of closed form local (or global) minima of L_N . A sufficient condition for B^\dagger to be nonnegative definite is that the distance involved be Euclidean (Borg and Groenen, 2005, Ch. 19). Although our metric ρ_s is not Euclidean, B^\dagger appears to be nonnegative definite in data analyses we have encountered, and the use of \mathbf{z}_\dagger has shown a good performance.

3.2 Proposed set classification procedure

The principal component features obtained from the previous section are now combined with the sample mean $\hat{\boldsymbol{\mu}}_i$, on which a classification rule is trained. Specifically, our set classification rule is learned in two steps:

Step 1: (Feature extraction)

- (a) For a given r , obtain the sample mean $\hat{\boldsymbol{\mu}}_i$ and PC spaces $\mathcal{L}_i^{(r)}$ for each set (Section 3.1.1);
- (b) Use CMDs to obtain the empirical PC space features \mathbf{z}_i (Sections 3.1.2 and 3.1.3);
- (c) Collect combined features $\mathbf{z}_i^* = [\hat{\boldsymbol{\mu}}_i; \mathbf{z}_i] \in \mathbb{R}^{p+m}$ ($i = 1, \dots, N$).

Step 2: Build a classification rule with any off-the-shelf methods, with inputs $(\mathbf{z}_i^*, y_i)_{i=1, \dots, N}$.

Prediction of the class label of a new observation $\mathcal{X}_\dagger = \{\mathbf{x}_1^\dagger, \dots, \mathbf{x}_{n_\dagger}^\dagger\}$ is also done in two steps. First obtain the features of the new set, $\mathbf{z}_\dagger^* = [\hat{\boldsymbol{\mu}}_\dagger; \mathbf{z}_\dagger]$, where $\hat{\boldsymbol{\mu}}_\dagger$ is the mean vector of the observation in the set \mathcal{X}_\dagger and \mathbf{z}_\dagger is from (4). The classification rule trained in Step 2 above is applied to the input \mathbf{z}_\dagger^* to predict the label \hat{y}_\dagger of the new set \mathcal{X}_\dagger .

4 Principal component space selection

The choice of an appropriate dimension r for the PC subspaces is a crucial step in feature extraction. We propose selecting an optimal dimension using training samples without making reference to the specific classification rules used in Step 2 above. A test to check whether one should choose $r = 0$, *i.e.*, discarding the PC features, is also considered.

Our strategy for selection of the dimension r of the PC spaces $\mathcal{L}_i^{(r)}$ is to use a modified Hotelling's T -squared statistic defined for each $1 \leq r \leq R$, where $R = \min\{n_1, \dots, n_N, p\}$ is the greatest meaningful dimension we can choose. Here, denote the coordinate matrix of $\{\mathcal{L}_i^{(r)}\}$ for each r by $\mathbf{Z}^{(r)} = (\mathbf{z}_i^{(r)})_{i=1}^N$, as defined in Theorem 2. Let G_1 and G_2 be a partition of $\{1, \dots, N\}$, where the set G_1 (or G_2) contains the indices corresponding to group 1 (or 2). Let $N_k = |G_k|$, $k = 1, 2$. Denote the mean difference by $\boldsymbol{\eta}_r = \bar{\mathbf{z}}_{G_1}^{(r)} - \bar{\mathbf{z}}_{G_2}^{(r)}$, where $\bar{\mathbf{z}}_{G_k}^{(r)} = \frac{1}{N_k} \sum_{i \in G_k} \mathbf{z}_i^{(r)}$, the pooled covariance matrix by $S_r = \frac{1}{N} \sum_{k=1}^2 \sum_{i \in G_k} (\mathbf{z}_i^{(r)} - \bar{\mathbf{z}}_{G_k}^{(r)})(\mathbf{z}_i^{(r)} - \bar{\mathbf{z}}_{G_k}^{(r)})^T$, and the diagonal matrix consisting of the diagonal elements of S_r by $D_r = \text{diag}(S_r)$.

The modified Hotelling's T^2 statistic for comparing two multivariate populations is the sum of squared marginal t -statistics $T(r) = \boldsymbol{\eta}_r' D_r^{-1} \boldsymbol{\eta}_r$, ($r = 1, \dots, R$) (Srivastava and Du, 2008). We choose an r that gives the greatest value of $T(r)$, *i.e.*,

$$\hat{r} = \hat{r}(T) = \underset{r=1, \dots, R}{\operatorname{argmax}} T(r), \quad (6)$$

since the greater $T(r)$ is, the more separable the groups are. A rationale for the diagonal matrix D_r in $T(r)$ comes from the use of the CMDS in Section 3.1.3. The coordinates of the features $\mathbf{z}_i^{(r)}$ are chosen so that they are in fact the principal axes; there is no correlation between coordinates of $\mathbf{Z}^{(r)}$. Moreover, with a small N , the inversion of the matrix S_r may cause numerical artifacts. Thus it is natural to exploit only the variances while discarding the covariances of $\mathbf{z}_i^{(r)}$. Accordingly, D_r is used instead of S_r .

A permutation test for usefulness of the chosen PC features $\{\mathbf{z}_i^{(\hat{r})}\}$ is now discussed. The features $\{\mathbf{z}_i^{(\hat{r})}\}$ may not contain any discriminating information when, for example, the

sub-populations differ only by location. This is translated into the null hypothesis:

$$H_0 : \mathbf{z}_i^{(\hat{r})}, i = 1, \dots, N, \text{ are identically distributed.}$$

The test statistic $T(\hat{r})$ tends to be large when the alternative hypothesis (not H_0) is true.

Since the group labels are exchangeable under the null hypothesis, the sampling distribution of $T(\hat{r})$ under H_0 is obtained by permuting the class labels. Let $\{T(r, b)\}_{r=1}^R$ be the re-calculated statistics using permuted data (the b th random permutation). The test then rejects H_0 if the test statistic $T(\hat{r})$ is larger than the $100(1 - \alpha)$ th percentile of the permuted statistics $T_b = \max_{1 \leq r \leq R} T(r, b)$ for a level of test α . In other words, the p -value of $T(\hat{r})$ is given by $p_{T(\hat{r})} = \frac{1}{B} \sum_{b=1}^B 1_{\{T(\hat{r}) \leq T_b\}}$ using B random permutations.

The dimension $\hat{r}(\geq 1)$ chosen by (6) is used for the PC features if $p_{T(\hat{r})} < \alpha$; otherwise we update $\hat{r} = 0$ so that the insignificant PC spaces are discarded.

When the true underlying model has no covariance difference between groups, the proposed permutation test greatly enhances the performance of our classification rules, as demonstrated in Web Appendix D.

5 Numerical studies

5.1 Competing methods

Our feature selection-set classification methods are denoted by *PCF-classifier*. For example, PCF-LDA denotes a set classification rule trained by linear discriminant analysis (LDA) on the extracted features, as that obtained in Sections 3 and 4. For the level of the permutation test, $\alpha = 0.05$ was used. Classifiers considered in numerical studies include LDA, quadratic discriminant analysis (QDA), and Support Vector Machines (SVM, Cristianini and Shawe-Taylor, 2000), Distance-Weighted Discrimination (DWD, Marron et al., 2007), and Minimum Distance Empirical Bayes rule (MDEB, Srivastava and Kubokawa, 2007).

Competing methods include voting classifiers, including two previous methods in Ning

and Karypis (2008), Wang et al. (2010). Consider binary classification with $Y = \pm 1$ and equal class sizes. Using LDA trained from all training samples, ignoring the set memberships, each \mathbf{x}_j in a new set $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ can be classified to $\hat{y}_j = \text{sign}\{(\mathbf{x}_j - \frac{\hat{\mu}_+ + \hat{\mu}_-}{2})^T \hat{\Sigma}^{-1}(\hat{\mu}_+ - \hat{\mu}_-)\}$. Wang et al. (2010) proposed to use a set classifier $\hat{y} = \sum_{j=1}^n \hat{y}_j$ determined by a majority vote of the individual predictions of \mathbf{x}_j , thus called an *LDA-MV* classifier. Weighted voting of the individual predictors, on the other hand, takes place when an *LDA-WV* classifier $\text{sign}\{\sum_{j=1}^n (\mathbf{x}_j - \frac{\hat{\mu}_+ + \hat{\mu}_-}{2})^T \hat{\Sigma}^{-1}(\hat{\mu}_+ - \hat{\mu}_-)\}$ is used. The LDA-WV classifier is an optimal classifier under the simple model (*cf.* Section 2). Classifiers QDA-(MV,WV), MDEB-(MV,WV), SVM-MV and DWD-MV are defined similar to LDA-(MV,WV) using the chosen classifiers. Precise definitions of these classifiers are provided in Web Appendix E. Ning and Karypis (2008) proposed using SVM-MV in classifications of sets.

In summary, the set classification methods we have considered are categorized into the following.

1. Proposed set classifiers: PCF-LDA, PCF-QDA, PCF-SVM, PCF-DWD and PCF-MDEB;
2. Existing methods: LDA-MV and SVM-MV;
3. Other competing methods: LDA-WV, QDA-(MV,WV), MDEB-(MV,WV) and DWD-MV.

5.2 Simulation models and results

We use the following four hierarchical models for binary set classification. Denote $\mathbf{e}_k \in \mathbb{R}^p$ for the vector whose components are all zeros except for a 1 in its k th. We have $\mathbf{x}_{ij} \sim N_p(\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)$ ($j = 1, \dots, n_i$), where $\boldsymbol{\mu}_i \mid (Y = k) \sim N_p(\boldsymbol{\delta}_{(k)}, 10^{-2} \mathbf{I}_p)$, $\boldsymbol{\delta}_{(1)} = \delta \mathbf{e}_1$, $\boldsymbol{\delta}_{(2)} = \mathbf{0}$, and

1. No covariance difference: $\boldsymbol{\Sigma}_i \equiv \boldsymbol{\Omega}_p(\rho)$;
2. Wishart: $\boldsymbol{\Sigma}_i \mid (Y = k) \sim \frac{1}{m} \mathcal{W}_p(\mathbf{V}_k, m)$, where $\mathbf{V}_k = \boldsymbol{\Omega}_p(\rho) + \sigma^2 \mathbf{e}_k \mathbf{e}_k'$ and $m = 10$;
3. Inverse Wishart: $\boldsymbol{\Sigma}_i \mid (Y = k) \sim p \mathcal{IW}_p(\mathbf{V}_k, p)$, where $\mathbf{V}_k = \boldsymbol{\Omega}_p(\rho) + \sigma^2 \mathbf{e}_k \mathbf{e}_k'$;
4. von Mises–Fisher: $\boldsymbol{\Sigma}_i = \boldsymbol{\Omega}_p(\rho) + \sigma^2 \mathbf{u}_i \mathbf{u}_i'$, where $\mathbf{u}_i \mid (Y = k) \sim \text{vMF}(\mathbf{e}_k, \kappa)$, and

$\kappa = 100$.

We use a class of covariance matrices $\mathbf{\Omega}_p(\rho)$ to consider highly correlated variables (Srivastava and Kubokawa, 2007, p.129). The $\mathbf{\Omega}_p(\rho)$ is a modified auto-regressive covariance matrix and is $\mathbf{\Omega}_p(\rho) = \left(\sigma_i \rho^{|i-j|^\frac{1}{p}} \sigma_j \right)_{ij}$, for $0 \leq \rho \leq 1$, where σ_i are independently drawn from the uniform distribution on $[4/5, 6/5]$. The von Mises–Fisher distribution used in model (4) is a Gaussian-like distribution for direction vectors, with mode at \mathbf{e}_k and a concentration parameter κ (cf. Mardia and Jupp, 2000). In all of the models, $\boldsymbol{\mu}_i$ and $\boldsymbol{\Sigma}_i$ are independent.

For each model, the performances of the methods in Section 5.1 are evaluated by empirical misclassification rates. The parameters of the models are set as $\delta = 1$, $\sigma = 3$, and $\rho \in \{0, .5\}$. We tested the number of sets $N = 10, 20$ (equal class sizes) with $p = 20, 200, 400$, and the set size n_i was independently selected by $n_i \sim \max\{\lfloor N(20, 5^2) \rfloor, 10\}$.

Table 1 summarizes the result of experiments based on 100 repetitions for $(p, N) = (20, 10), (400, 20)$. Throughout different models and various settings of dimension–sample size pairs, the proposed methods exhibit much smaller misclassification rates than the competing methods. This is not surprising because the models (2–4) from which the data are generated are the hierarchical models with a difference in PC structures. Since our method choose the PC spaces as features, good performance is to be expected. Moreover, even when there is no covariance difference in model (1), the proposed methods are comparable to the weighted voting methods. In model (1) with $\rho = 0$, the permutation test in Section 4 makes the PC space features used only 2–8% of the time, thus effectively discarding unimportant PC information. On the other hand, in models (2–4) where the PC spaces possess discriminating information related to class, the tests resulted in rejecting the null hypothesis with empirical power of 95–100%, thus successfully utilizing the PC space features. When the variables are highly correlated ($\rho = 0.5$), our methods show rates smaller or comparable to those obtained using MDEB-WV, which is designed to work well for correlated variables. The simulation result also confirms that LDA-WV is not optimal in the general hierarchical model.

Table 1: Empirical misclassification rates (in %, mean and standard error based on 100 simulations) illustrate that our methods PCF-‘classifier’ show better performance than competing methods. Full tables including the results of all dimension-sample size pairs and all methods listed in Section 5.1 are available at Appendix Table F.1 and F.2.

Model	(p, N)	Proposed methods (<i>PCF-‘classifier’</i>)						Competing methods		
		PCF-LDA	PCF-SVM	PCF-DWD	PCF-MDEB	LDA-WV	SVM-MV	MDEB-WV		
(1), $\rho = 0$	(20,10)	25.41(5.43)	22.03(5.45)	18.75(4.28)	18.26(4.42)	18.18(4.60)	20.72(4.56)	19.57(6.08)		
	(400,20)	33.18(3.86)	33.22(3.89)	32.61(3.97)	32.59(3.93)	37.61(4.47)	38.27(4.70)	33.42(4.01)		
(2), $\rho = 0$	(20,10)	4.06(3.58)	5.20(4.39)	6.56(4.76)	6.14(4.60)	38.18(4.72)	39.95(4.80)	34.34(7.26)		
	(400,20)	2.22(1.25)	2.22(1.25)	2.22(1.24)	2.22(1.26)	47.80(3.66)	46.55(3.37)	43.31(3.31)		
(3), $\rho = 0$	(20,10)	5.80(4.07)	5.75(3.83)	6.10(3.79)	7.34(5.69)	36.10(6.33)	36.74(6.14)	34.81(12.20)		
	(400,20)	21.03(6.94)	21.03(6.94)	20.88(6.74)	23.35(7.18)	34.67(3.48)	32.57(3.66)	30.14(6.78)		
(4), $\rho = 0$	(20,10)	13.78(5.66)	15.53(5.41)	15.73(4.81)	15.36(4.55)	37.25(4.36)	38.28(4.20)	33.83(6.30)		
	(400,20)	32.49(3.58)	32.50(3.55)	31.83(3.56)	31.65(3.53)	38.24(4.94)	38.72(4.39)	33.55(3.92)		
(1), $\rho = .5$	(20,10)	23.69(5.49)	20.57(4.96)	17.44(3.96)	17.09(3.99)	16.56(4.07)	18.39(3.76)	17.31(5.24)		
	(400,20)	33.93(4.75)	34.01(4.81)	34.28(4.76)	34.41(4.82)	36.13(4.51)	36.08(4.12)	31.39(3.96)		
(2), $\rho = .5$	(20,10)	3.91(4.23)	4.68(4.55)	6.16(4.95)	5.71(4.81)	39.52(4.66)	40.25(4.94)	34.06(7.11)		
	(400,20)	1.58(1.05)	1.58(1.05)	1.56(1.05)	1.58(1.05)	48.24(3.46)	47.00(3.31)	43.45(3.35)		
(3), $\rho = .5$	(20,10)	5.55(4.46)	5.46(4.14)	5.96(4.26)	7.38(7.21)	37.28(6.41)	37.96(6.37)	34.66(11.99)		
	(400,20)	20.40(7.09)	20.46(7.13)	20.19(6.96)	22.38(6.07)	35.19(3.54)	32.98(3.81)	30.16(6.70)		
(4), $\rho = .5$	(20,10)	1.97(1.70)	2.86(2.10)	4.46(2.77)	3.76(2.29)	38.01(3.96)	38.46(4.31)	32.34(6.18)		
	(400,20)	33.32(4.42)	33.34(4.40)	33.89(4.88)	33.95(5.01)	37.83(4.73)	37.48(4.40)	31.80(3.52)		

6 Classification of liver cell nucleus images

One of the common procedures used to diagnose hepatoblastoma (a rare malignant liver cancer) is biopsy—a sample tissue of a tumor is removed and examined under a microscope. A tissue sample contains a number of nuclei, a subset of which is then processed to obtain segmented images of nuclei. The data we analyzed contain 5 sets of nuclei from normal liver tissues and 5 sets of nuclei from cancerous tissues; see Figure 1. Each set contains 50 images. The data set is available at <http://www.andrew.cmu.edu/user/gustavor/software.html> and is introduced in Wang et al. (2010, 2011), in which one can find details on registration and segmentation of the nucleus images.

We tested the proposed method on the liver cell nuclei image data set in two different ways. First, using all 50 observations in each set, we evaluated the leave-one-out classification error rates by fixing one set as test data and training with the remaining 9 sets. This experiment was repeated for dimension-reduced data (using PCA for combined data). Table 2 summarizes the results. Our methods only misclassified the 9th set, showing the best performance (together with LDA-MV) among all methods considered. The 9th set was misclassified by many methods in different dimensions. This can be explained by a visual inspection of the scatterplot in Figure 2. The orientation of the distribution of the 9th set, illustrated as the thicker orange ellipse in the right panel, is pointing north-east, which is closer to the orientation of the normal group than the north-west orientation of the hepatoblastoma group.

Second, the performances of the set classifiers were evaluated when the number of observations n in each set is smaller than 50. Although technology enables processing of an image through semi-automatic registration and segmentation, the cost of obtaining images is still expensive. Thus a method exhibiting solid performance in small n setting has a clear advantage over other methods. In particular, we randomly chose $n \in \{5, 10, 15, 20, 25\}$ images in each of $N = 10$ sets as training data and also chose non-overlapping n images in each set as testing data. After training the set classifiers, the empirical misclassification

Table 2: Misclassification rates and labels of misclassified sets in leave-one-out classification experiments of liver data. From original data with dimension $d = 22,500$ to dimension-reduced data by PCA. Results from QDA-(MV, WV) and DWD-MV are omitted as their performances are poor.

	$d = 36,864$	$d = 100$	$d = 10$	$d = 2$
PCF-LDA	1/10, (9th)	1/10, (9th)	0/10	0/10
PCF-QDA	1/10, (9th)	1/10, (9th)	0/10	2/10, (4,5)th
PCF-SVM	1/10, (9th)	1/10, (9th)	0/10	0/10
PCF-DWD	1/10, (9th)	1/10, (9th)	0/10	0/10
PCF-MDEB	1/10, (9th)	1/10, (9th)	0/10	0/10
LDA-MV	1/10, (9th)	1/10, (9th)	1/10, (9th)	1/10, (9th)
LDA-WV	3/10, (3,7,9)th	0/10	1/10, (9th)	1/10, (9th)
SVM-MV	1/10, (9th)	2/10, (6,9)th	1/10, (9th)	1/10, (9th)
MDEB-MV	1/10, (9th)	1/10, (9th)	1/10, (9th)	3/10, (3,7,9)th
MDEB-WV	1/10, (9th)	1/10, (9th)	0/10	2/10, (7,9)th

rate was computed using the remaining testing data. This procedure was repeated for 100 random subsets of size n , for each choice of n . Figure 4 summarizes the results of the small- n experiment. Our methods exhibited significantly better performance than the others, with the exception of MDEB-WV, for both the original and dimension-reduced data sets.

The features of the data we chose only include ‘statistical features’ of the raw images, but do not include any other features that needs expert input, such as shape or texture features (*cf.* [Chen et al., 2006](#)).

7 Discussion

We have introduced a novel set classification framework and proposed a feature extraction–selection procedure. The method is based on the hierarchical modeling of set classification data. The features of sets, chosen by a data-driven approach, are used as inputs for any off-the-shelf classification method. When the orientation of major variations in sets possesses discriminating information, the proposed method adaptively uses such information for classification. Classical multidimensional scaling with the modified Euclidean sine metric between PC subspaces are shown to work well in mapping nonlinear features into vector-valued




Figure 4: Misclassification rates by different set size n for the binary set classification of liver cell nuclei images. The proposed set classifiers exhibit much smaller misclassification rates than other methods when both n and N (the number of sets) are small. The weighted voting classifier with MDEB rule (MDEB-WV) shows comparable performance. An initial dimension reduction (PCA for combined data) improves overall performances of most classification methods.

points. An analytic solution for the mapping of new observations is derived.

Note that [Cox and Ferry \(1993\)](#) and [Witten and Tibshirani \(2011\)](#) also proposed using multidimensional scaling (MDS) for classification. However, the method proposed in Cox and Ferry only works for vector inputs and uses a linear regression for mapping of new observations. Witten and Tibshirani proposed modifying the nonmetric multidimensional scaling, which requires an iterative algorithm to solve, but they did not consider using MDS for set classification.

There is the potential to extend our framework to more general cases. In particular, the independence assumption in the hierarchical model may be relaxed to model dependencies among observations in a set. This is especially relevant in classification of tissues based on sets of images because the cells in a tissue have location information. Images observed with close locations can be assumed to have high correlation. We conjecture that including the location in the hierarchical model greatly increases the accuracy of classification. We have focused on the independent hierarchical model in this paper, as it serves well as an introduction to set classification.

The performance of our set classifiers can be greatly limited when the set sizes are severely imbalanced. The principal component estimates for each set use only n_i samples. If n_i vary greatly, so does the credibility of the estimated features. Moreover, the choice of the dimension r for the PC features in [Section 4](#) is also limited up to $\min(n_i)$. These problems are mitigated by excluding those sets with a small sample size or by applying a sensible dimension reduction, but the classifiers could be further honed to properly accommodate imbalanced set sizes.

Another issue is the use of classifiers. We have demonstrated the use of conventional classifiers (LDA, QDA and linear SVM), but have not examined using more sophisticated methods such as kernel machines (*cf.* [Cristianini and Shawe-Taylor, 2000](#)). While use of kernel machines may benefit the classification task as shown in [Wang et al. \(2010\)](#), we chose to focus on our feature selection procedure using simpler methods. On the other hand, since

our method is a general framework that works with any classifier, the presented methods have the potential to be improved by advanced kernel machines.

8 Supplementary Materials

Web Appendices referenced in Sections 2, 3.1, 4 and 5 as well as a zipped code and example data are available.

Acknowledgements

We thank Gustavo Rohde for discussions on related topics and providing the histopathology data sets, and the Editor, Associate Editor and anonymous referees for their valuable comments. Jung was partially supported by NSF grant DMS-1307178. Qiao was partially supported by the Simons Foundation #246649.

References

- Ahn, J. and Marron, J. S. (2010), “The maximal data piling direction for discrimination,” *Biometrika*, 97, 254–259.
- Ahn, J., Marron, J. S., Muller, K. M., and Chi, Y.-Y. (2007), “The high-dimension, low-sample-size geometric representation holds under mild conditions,” *Biometrika*, 94, 760–766.
- Bickel, P. and Levina, E. (2004), “Some theory for Fisher’s linear discriminant function, ‘naive Bayes’, and some alternatives when there are many more variables than observations,” *Bernoulli*, 10, 989–1010.
- Borg, I. and Groenen, P. J. (2005), *Modern Multidimensional Scaling: Theory and Applications*, Springer Series in Statistics, 2nd ed.

- Bradley, R. C. (2005), “Basic properties of strong mixing conditions. A survey and some open questions,” *Probab. Surv.*, 2, 107–144 (electronic), update of, and a supplement to, the 1986 original.
- Casella, G. and Berger, R. L. (2002), *Statistical inference*, Pacific Grove, CA: Duxbury, 2nd ed.
- Chen, Y. C., Hu, K. H., Li, F. Z., Li, S. Y., Su, W. F., Huang, Z. Y., and Hu, Y. X. (2006), “Studies on quantitative analysis and automatic recognition of cell types of lung cancer,” *Biomed. Mater. Eng.*, 16, 119–128.
- Chikuse, Y. (2003), *Statistics on special manifolds*, New York: Springer-Verlag.
- Cox, T. F. and Cox, M. a. a. (2008), “Multidimensional Scaling,” in *Springer Handbooks of Computational Statistics 2008, Handbook of Data Visualization, III*, eds. Chen, C.-h., Härdle, W., and Unwin, A., Springer, vol. 38, pp. 315–347.
- Cox, T. F. and Ferry, G. (1993), “Discriminant analysis using non-metric multidimensional scaling,” *Pattern Recog.*, 26, 145–153.
- Cristianini, N. and Shawe-Taylor, J. (2000), *An Introduction to Support Vector Machines*, Cambridge University Press.
- Dempster, A. P. (1960), “A significance test for the separation of two highly multivariate small samples,” *Biometrics*, 16, 41–50.
- Friedman, J. H. (1989), “Regularized discriminant analysis,” *J. Amer. Statist. Assoc.*, 84, 165–175.
- Gower, J. C. (1966), “Some Distance Properties of Latent Root and Vector Methods Used in Multivariate Analysis,” *Biometrika*, 53, 325–338.

- Hall, P., Marron, J. S., and Neeman, A. (2005), “Geometric representation of high dimension, low sample size data,” *Journal of the Royal Statistical Society - Series B: Statistical Methodology*, 67, 427–444.
- Jolliffe, I. T. (2002), *Principal Component Analysis*, New York: Springer-Verlag, 2nd ed.
- Jung, S. and Marron, J. S. (2009), “PCA consistency in high dimension, low sample size context,” *The Annals of Statistics*, 37, 4104–4130.
- Jung, S., Sen, A., and Marron, J. S. (2012), “Boundary behavior in high dimension, low sample size asymptotics of PCA,” *The Journal of Multivariate Analysis*, 109, 190–203.
- Lee, M. H. (2007), “Continuum Direction Vectors in High Dimensional Low Sample Size Data,” Ph.D. thesis, University of North Carolina at Chapel Hill.
- Mardia, K. V. and Jupp, P. E. (2000), *Directional statistics*, Chichester: John Wiley & Sons Ltd.
- Marron, J. S., Todd, M. J., and Ahn, J. (2007), “Distance Weighted Discrimination,” *J. Amer. Statist. Assoc.*, 102, 1267–1271.
- Nie, J. (2006), “Global Optimization of Polynomial Functions and Applications,” Ph.D. thesis, University of California, Berkeley.
- Ning, X. and Karypis, G. (2008), “The Set Classification Problem and Solution Methods,” in *IEEE International Conference on Data Mining Workshops*, pp. 720–729.
- Qiao, X., Zhang, H. H., Liu, Y., Todd, M. J., and Marron, J. S. (2010), “Weighted distance weighted discrimination and its asymptotic properties,” *J. Amer. Statist. Assoc.*, 105, 401–414.
- Samsudin, N. A. and Bradley, A. P. (2010), “Nearest neighbour group-based classification,” *Pattern Recogn.*, 43, 3458–3467.

- Srivastava, M. S. and Du, M. (2008), “A test for the mean vector with fewer observations than the dimension,” *J. Multivariate Anal.*, 99, 386–402.
- Srivastava, M. S. and Fujikoshi, Y. (2006), “Multivariate analysis of variance with fewer observations than the dimension,” *J. Multivariate Anal.*, 97, 1927–1940.
- Srivastava, M. S. and Kubokawa, T. (2007), “Comparison of discrimination methods for high dimensional data,” *J. Japan Statist. Soc.*, 37, 123–134.
- Stewart, G. W. and Sun, J. G. (1990), *Matrix perturbation theory*, Boston: Academic Press.
- Tsantis, S., Dimitropoulos, N., Cavouras, D., and Nikiforidis, G. (2009), “Morphological and wavelet features towards sonographic thyroid nodules evaluation,” *Comput. Med. Imag. Grap.*, 33, 91–99.
- Wang, W., Ozolek, J. A., and Rohde, G. K. (2010), “Detection and classification of thyroid follicular lesions based on nuclear structure from histopathology images,” *Cytom. Part A*, 77, 485–494.
- Wang, W., Ozolek, J. A., Slepčev, D., Lee, A. B., Chen, C., and Rohde, G. K. (2011), “An optimal transportation approach for nuclear structure-based pathology,” *IEEE Trans. Med. Imaging*, 30, 621631.
- Witten, D. M. and Tibshirani, R. (2011), “Supervised multidimensional scaling for visualization, classification, and bipartite ranking,” *Comput. Stat. Data Anal.*, 55, 789 – 801.
- Yata, K. and Aoshima, M. (2012), “Effective PCA for high-dimension, low-sample-size data with noise reduction via geometric representations,” *J. Multivariate Anal.*, 105, 193–215.

Web-based Supplementary Materials: “A Statistical Approach to Set Classification by Feature Selection with Applications to Classification of Histopathology images”

by Sungkyu Jung and Xingye Qiao

University of Pittsburgh and State University of New York, Binghamton

Summary

This supplementary document for ‘A Statistical Approach to Set Classification by Feature Selection with Applications to Classification of Histopathology images’ includes

1. Optimal decision rules in set classification as referenced in Section 2 (Web Appendix [A.1–A.3](#));
2. A proof of Proposition 1 as referenced in Section 3.1.1 (Web Appendix [A.4](#));
3. A proof of Theorem 2 as referenced in Section 3.1.3 (Web Appendix [A.5](#));
4. A discussion on the scale factor as referenced in Section 3.1.2 (Web Appendix [B](#));
5. Additional tables and figures as referenced in Section 4 (Web Appendix [C](#) and [D](#));
6. Additional definitions and tables as referenced in Section 5 (Web Appendix [E](#) and [F](#)).
7. A collection of Matlab codes for the proposed set classification and example data are available as a zipped file.

A Technical details

In this section, we first provide theoretical optimal classifiers for the set classification models introduced in Section 2, and discuss their implications to the proposed method. In sections [A.4](#) and [A.5](#) proofs of Proposition 1 and Theorem 3 are provided.

A.1 Simple model

The simple model for set classification, referenced in Section 2, is a special case of the hierarchical model and can be written as

$$\mathcal{X}_i \mid (Y = k) \sim \prod_{j=1}^{n_i} f_k(\mathbf{x}_{ij})$$

with $P(\{Y = k\}) = \pi_k$.

Let (\mathcal{X}, Y) be a new random object from the simple model. Let Φ be a collection of set classification rules, and $\phi(\mathcal{X}) \equiv \phi(\{\mathbf{x}_1, \dots, \mathbf{x}_n\}) \in \Phi$. The theoretical Bayes optimal decision rule ϕ^* minimizes the risk function $R(\phi) = P(\phi(\mathcal{X}) \neq Y)$ among all $\phi \in \Phi$. Since

$$\begin{aligned} 1 - R(\phi) &= \sum_{k=1}^K \pi_k P(\phi(\mathcal{X}) = k \mid Y = k) \\ &= \sum_k \pi_k \int \mathbb{1}_{\{\phi(\mathcal{X})=k\}} f_k(\mathcal{X}) d\mathcal{X}, \end{aligned}$$

the optimal ϕ^* must satisfy $\phi^*(\mathcal{X}) = \operatorname{argmax}_k f_k(\mathcal{X})\pi_k = \operatorname{argmax}_k \prod_{j=1}^n f_k(\mathbf{x}_j)\pi_k$ for any \mathcal{X} .

With an assumption that each f_k is a p -variate normal distribution with mean $\boldsymbol{\mu}_k$ and covariance matrix Σ , a classifier using the linear discriminant analysis can be shown to be optimal in the following sense. For simplicity, consider binary classification with $Y = \pm 1$ and assume $\pi_+ = \pi_-$. Then the Bayes decision rule is

$$\begin{aligned} \phi^*(\mathcal{X}) &= \operatorname{argmax}_{k \in \{+1, -1\}} \prod_{j=1}^n f_k(\mathbf{x}_j) \\ &= \operatorname{sign} \left\{ \left(\prod_{j=1}^n f_+(\mathbf{x}_j) / \prod_{j=1}^n f_-(\mathbf{x}_j) \right) - 1 \right\} \\ &= \operatorname{sign} \left\{ \sum_{j=1}^n \left(\mathbf{x}_j - \frac{\boldsymbol{\mu}_+ + \boldsymbol{\mu}_-}{2} \right)^T \Sigma^{-1} (\boldsymbol{\mu}_+ - \boldsymbol{\mu}_-) \right\}. \end{aligned} \tag{7}$$

An estimator $\hat{\phi}^*(\mathcal{X})$ of the optimal rule can be obtained by plugging-in the estimates $\hat{\boldsymbol{\mu}}_+, \hat{\boldsymbol{\mu}}_-$ and $\hat{\Sigma}$. If the order of the sign operator and summation in (7) is switched, then a related,

but suboptimal, classifier is obtained. The corresponding empirical classifier

$$\sum_{j=1}^n \text{sign} \left\{ \left(\mathbf{x}_j - \frac{\hat{\mu}_+ + \hat{\mu}_-}{2} \right)^T \hat{\Sigma}^{-1} (\hat{\mu}_+ - \hat{\mu}_1) \right\} \quad (8)$$

works as if the set label is determined by a majority vote of the individual predictions of \mathbf{x}_j . We call (8) as an *LDA-MV* classifier. In contrast, the optimal classifier (7) works as if the individual predictions vote with weights, thus called a *LDA-WV* (weighted voting) classifier.

The LDA-WV is optimal only if the simple model can be assumed. In the hierarchical model, both LDA-MV and LDA-WV exhibit poor performances in the simulated and real data analyses, shown in Section 5 and 6. The majority voting strategy has been suggested in [Ning and Karypis \(2008\)](#) and [Wang et al. \(2010\)](#), but they have not discussed its relationship to the theoretical Bayes rule.

A.2 Hierarchical model

The hierarchical model (equation 1, Section 2) enables us to model appropriate characteristics of set classification data. In the hierarchical model, the optimal classification rule for prediction of a new set \mathcal{X} minimizes the risk $P(\phi(\mathcal{X}) \neq Y)$. Writing $\pi_k = P(\{Y = k\})$, we have

$$\begin{aligned} P(\phi(\mathcal{X}) = Y) &= \sum_{k=1}^K \pi_k P(\phi(\mathcal{X}) = k \mid Y = k) \\ &= \sum_{k=1}^K \pi_k \mathbb{E}_{\Theta} [P(\phi(\mathcal{X}) = k \mid \Theta) \mid Y = k] \\ &= \sum_{k=1}^K \int \mathbb{1}_{\{\phi(\mathcal{X})=k\}} \left(\pi_k \int f(\mathcal{X}; \theta) h_k(\theta) d\theta \right) d\mathcal{X}, \end{aligned} \quad (9)$$

where $h_k(\theta)$ is the conditional density function of Θ given $Y = k$, provided that the density exists. From (9), the optimal rule $\phi(\mathcal{X})$ predicts $\hat{Y} = k$ if k maximizes

$$\pi_k \int f(\mathcal{X}; \theta) h_k(\theta) d\theta. \quad (10)$$

Since the estimation of (10) is challenging, we look at the problem at a different angle.

First consider the special case that the parameter Θ is fixed for each class k , *i.e.*, $P(\Theta = \theta \mid Y = k) = \mathbb{1}_{\{\theta = \theta_k\}}$. Then (10) is simplified to $\pi_k f(\mathcal{X}; \theta_k)$, which is the expression used in the optimal classifier of the simple model. On the other hand, if the distribution of \mathcal{X} given θ is concentrated at only one point, *i.e.*, $P(\mathcal{X} = X \mid \Theta = \theta) = \mathbb{1}_{\{g(X) = \theta\}}$ for an invertible function g , then the optimal classifier is

$$\phi(\mathcal{X}) = \operatorname{argmax}_{k=1,\dots,K} \pi_k h_k(g(\mathcal{X})) = \operatorname{argmax}_{k=1,\dots,K} \pi_k h_k(\theta_{(\mathcal{X})}). \quad (11)$$

In the latter case, a classifier based on \mathcal{X} is as good as a classifier based on the parameter $\theta_{(\mathcal{X})}$ of the set. This theoretical observation, as well as the empirical observation from the real data in Fig. 2, has contributed in considering a classification based on the parameter θ (or the features of a set) rather than individual observations in \mathcal{X} .

The proposed set classification procedure can be thought of as building a classification rule on the predicted parameter (feature) $\hat{\theta}$. The parameter is modeled to include the mean and PC spaces. A remaining question is the quality of the prediction of θ . If $\hat{\theta}$ is close to θ , then the empirical classifier is as good as the optimal classifier (11). In the next section we show that the empirical PC space features are as good as the true PC spaces in the high dimensional asymptotic context.

A.3 High dimensional asymptotic theory

The quality of the empirical features $\mathcal{L}_i^{(r)}$ is now assessed in the high-dimensional, low-sample size (HDLSS) asymptotic context to support the theoretical motivation in Section A.2. Our image dataset suits well in HDLSS context since it is in very high dimension with p being tens of thousands while the number of sets N and the number of observations n in the set are both small. An asymptotic investigation where $p \rightarrow \infty$ and both (N, n) fixed is thus relevant to the analysis of such HDLSS data (Hall et al., 2005, Ahn et al., 2007).

Since we are interested in the empirical PC spaces, assume for $X_{ij(p)}$, the j th random observation in the i th set, $\mathbb{E}(X_{ij(p)}) = \mathbf{0}_p$ and $\operatorname{Cov}(X_{ij(p)}) = \Sigma_{i(p)} = \sum_{l=1}^p \lambda_{li(p)}^o e_{li(p)}^o (e_{li(p)}^o)^T$.

Here, the subscript (p) is used to emphasize the dependence on the dimension p . Let $\mathcal{U}_{i(p)}^{(r)} = \text{span}\{e_{1i(p)}^o, \dots, e_{ri(p)}^o\}$ be the true PC space, which is the target of the empirical PC space $\mathcal{L}_i^{(r)} \equiv \mathcal{L}_{i(p)}^{(r)}$.

The following theorem states that when the first r eigenvalues are much larger than the others, then the set classification based on the empirical PC spaces $\mathcal{L}_{i(p)}^{(r)}$ is as good as using the true PC spaces $\mathcal{U}_{i(p)}^{(r)}$ in the limit.

For two sequences a_p and b_p let $a_p \asymp b_p$ stand for the relation $\lim_{p \rightarrow \infty} a_p/b_p \in (0, \infty)$ and $\lim_{p \rightarrow \infty} a_p/b_p \in (0, \infty)$.

We need the following moment condition:

(C) Each components of $X_{ij(p)}$ have finite fourth moments for all p , and there exists a permutation so that the permuted sequence of loadings of the scaled PCs $Z_{ij(p)} = ((\lambda_{li(p)}^o)^{-\frac{1}{2}}(e_{li(p)}^o)^T X_{ij(p)})_{l=1, \dots, p}$ are ρ -mixing (Bradley, 2005).

THEOREM 4. *Assume that $\lambda_{li(p)}^o \asymp p^\alpha$, for $l = 1, \dots, r_0$ and the rest of eigenvalues are fixed, for example $\lambda_{li(p)}^o \equiv 1$ for $l = r_0 + 1, \dots, p$. Suppose n is fixed and (C) holds.*

- (a) *If $\alpha > 1$, then $\mathcal{L}_{i(p)}^{(r)}$ is consistent with $\mathcal{U}_{i(p)}^{(r)}$ in the sense that $\rho_s(\mathcal{L}_{i(p)}^{(r)}, \mathcal{U}_{i(p)}^{(r)}) \rightarrow 0$ in probability as $p \rightarrow \infty$.*
- (b) *If $\alpha < 1$, then $\mathcal{L}_{i(p)}^{(r)}$ is strongly inconsistent with $\mathcal{U}_{i(p)}^{(r)}$ in the limit $p \rightarrow \infty$ in the sense that $\rho_s(\mathcal{L}_{i(p)}^{(r)}, \mathcal{U}_{i(p)}^{(r)}) \rightarrow \sqrt{r}$ in probability as $p \rightarrow \infty$.*
- (c) *If $\alpha = 1$, $\rho_s(\mathcal{L}_{i(p)}^{(r)}, \mathcal{U}_{i(p)}^{(r)})$ weakly converges to a distribution with support on $[0, \sqrt{r}]$ as $p \rightarrow \infty$.*

Theorem 4 tells us that, in the limit $p \rightarrow \infty$, $\mathcal{L}_{i(p)}^{(r)}$ can be used in place of $\mathcal{U}_{i(p)}^{(r)}$ when signal from the PCs is strong (i.e., large $\lambda_{li(p)}^o$). Theorem 4 can be shown by an application of Theorem 2 of Jung and Marron (2009) (for $\alpha \neq 1$) and Theorem 3 of Jung et al. (2012) (for $\alpha = 1$) with the fact that $\sup_{U_r, V_r} \rho(U_r, V_r) = \sqrt{r}$ for $U_r, V_r \in G(r, p)$.

Next, we show that the classification based on $\mathcal{U}_{i(p)}^{(r)}$ is optimal in a simplified setting. Suppose that in the binary classification we have;

ASSUMPTION 1. $\text{Cov}(X_{ij(p)}) = \Sigma_{i(p)}$ is modeled with a single factor as $\Sigma_{i(p)} = \mathbf{I}_p + \mathbf{u}_{i(p)}\mathbf{u}_{i(p)}^T$, where $\mathbf{u}_{i(p)} \mid (Y = +1) \sim N_p(\boldsymbol{\mu}_+(p), \mathbf{I}_p)$ and $\mathbf{u}_{i(p)} \mid (Y = -1) \sim N_p(\boldsymbol{\mu}_-(p), \mathbf{I}_p)$.

The following theorem is the result of directly applying the geometric representation theory in [Hall et al. \(2005\)](#), [Ahn et al. \(2007\)](#) and [Qiao et al. \(2010\)](#). Note that we have assumed normality for $\mathbf{u}_{i(p)}$ to simplify the situation.

PROPOSITION 5. *Under Assumption 1, the pairwise distances between the N_+ (N_- , resp.) leading eigenvectors for sets from the positive (negative, resp.) class are approximately the same. In particular, scaled by $1/p$, the squared distances satisfy $\frac{1}{p}\|\mathbf{u}_k^+ - \mathbf{u}_l^+\|^2 \rightarrow 2$ and $\frac{1}{p}\|\mathbf{u}_k^- - \mathbf{u}_l^-\|^2 \rightarrow 2$.*

Now if we further assume that the squared mean difference for the \mathbf{u} 's is $\frac{1}{p^\alpha}\|\boldsymbol{\mu}_+(p) - \boldsymbol{\mu}_-(p)\|^2 \rightarrow \mu^2$, then we can derive the following asymptotic result which outlines the conditions under which the SVM classifier can always correctly classify the set data based on the leading eigenvectors.

THEOREM 6. *Without loss of generality, assume that $N_+ \leq N_-$, and hence $1/N_+ \geq 1/N_-$.*

1. *If $\alpha > 1$, then with probability converging to 1 as $p \rightarrow \infty$, a new set from either class will be correctly classified by the SVM classifier.*
2. *If $\alpha < 1$, then with probability converging to 1 as $p \rightarrow \infty$, a new set from either class will be misclassified by the SVM classifier.*
3. *If $\alpha = 1$, then when $\mu^2 > 1/N_+ - 1/N_-$, with probability converging to 1 as $p \rightarrow \infty$, a new set from either class will be correctly classified by the SVM classifier.*

A.4 Proof of Proposition 1

The identifiability of the parameter $\mathcal{L}^o = \text{span}(\mathbf{e}_1^o, \mathbf{e}_2^o)$ is assured by the assumption $\lambda_1^o = \lambda_2^o > \lambda_3^o$ and the fact that the spectral decomposition is unique up to sign change. See, for example, [Casella and Berger \(2002, p. 523\)](#) for the notion of identifiability.

The consistency of $\mathcal{L}_i^{(2)} = \text{span}(\mathbf{e}_{1i}, \mathbf{e}_{2i})$ with \mathcal{L}^o for the high-dimensional, low-sample size context ($p \rightarrow \infty$ with fixed n) is stated in Theorem 4. The large sample asymptotic consistency is trivially obtained again by the uniqueness of the spectral decomposition as well as the consistency of the empirical covariance matrix $\widehat{\Sigma}_i$ with Σ_i .

A.5 Proof of Theorem 3

With the notation $\mathbf{z} = (\mathbf{z}_{(1)}^T; z_{(2)})^T \in \mathbb{R}^{m+1}$, the loss function L_N is written as

$$\begin{aligned} L_N(\mathbf{z}) &= \left\| -\frac{1}{2}P_N[D(\mathbf{Z}_{\dagger}(\mathbf{z})) - \Delta_{\dagger}]P_N^T \right\|_F^2 = \|(\mathbf{Z}_{\dagger}(\mathbf{z}))^T \mathbf{Z}_{\dagger} - B^{\dagger}\|_F^2 \\ &= \|\mathbf{Z}^T \mathbf{Z} - B\|_F^2 + 2\|\mathbf{Z}^T \mathbf{z}_{(1)} - \mathbf{b}_{12}\|_F^2 + a(\mathbf{z})^2, \end{aligned}$$

where $a(\mathbf{z}) = a(\mathbf{z}_{(1)}, z_{(2)}) = \mathbf{z}_{(1)}^T \mathbf{z}_{(1)} + z_{(2)}^2 - b_2$. The first order condition leads that

$$\frac{1}{4} \frac{dL_N(\mathbf{z})}{d\mathbf{z}} = \frac{1}{4} \begin{pmatrix} \frac{dL_N(\mathbf{z})}{d\mathbf{z}_{(1)}} \\ \frac{\partial L_N(\mathbf{z})}{\partial z_{(2)}} \end{pmatrix} = \begin{bmatrix} \mathbf{Z}\mathbf{Z}^T \mathbf{z}_{(1)} - \mathbf{Z}\mathbf{b}_{12} + a(\mathbf{z})\mathbf{z}_{(1)} \\ 2(\mathbf{z}_{(1)}^T \mathbf{z}_{(1)} + z_{(2)}^2 - b_2)z_{(2)} \end{bmatrix} = \mathbf{0}. \quad (12)$$

The Hessian matrix is

$$\begin{aligned} H(\mathbf{z}) &= \begin{bmatrix} \mathbf{Z}\mathbf{Z}^T + \mathbf{z}_{(1)}\mathbf{z}_{(1)}^T + a(\mathbf{z})\mathbf{I}_m & 2z_{(2)}\mathbf{z}_{(1)} \\ 2z_{(2)}\mathbf{z}_{(1)}^T & 2a(\mathbf{z}) + 4z_{(2)}^2 \end{bmatrix} \\ &= \begin{bmatrix} \Lambda_+ + a(\mathbf{z})\mathbf{I}_m & \mathbf{0} \\ \mathbf{0}^T & 2a(\mathbf{z}) \end{bmatrix} + \begin{pmatrix} \mathbf{z}_{(1)} \\ 2z_{(2)} \end{pmatrix} \begin{pmatrix} \mathbf{z}_{(1)} \\ 2z_{(2)} \end{pmatrix}^T, \end{aligned}$$

since $\mathbf{Z}\mathbf{Z}^T = \Lambda_+$. Now the $\hat{\mathbf{z}}$ (defined in Theorem 2) satisfies (12), and $H(\hat{\mathbf{z}})$ is positive definite if $\hat{z}_{(2)}^2 = b_2 - \mathbf{z}_{\dagger}^T \mathbf{z}_{\dagger} > 0$, because $a(\hat{\mathbf{z}}) = 0$. If $\hat{z}_{(2)}^2 = 0$, then $L_N(\mathbf{z})$ is invariant to $z_{(2)}$, and $\hat{\mathbf{z}}$ is in fact the global minimum. This shows that $\hat{\mathbf{z}}$ are the local minima of L_N .

For the second argument, if it is nonnegative definite, B^{\dagger} can be thought of as a covariance matrix of an $(N+1)$ -variate normal distribution. We have $b_2 - \mathbf{z}_{\dagger}^T \mathbf{z}_{\dagger} = b_{22} - \mathbf{b}_{21}B^{-1}\mathbf{b}_{12} \geq 0$, since the latter is the conditional variance of the last coordinate given the first N coordinates, which is nonnegative.

B On the scale factor c of the Euclidean Sine metric (Section 3.1.2)

See Fig. 5.




Figure 5: The choice of c as the average of the empirical total variance in $\mathcal{L}_i^{(r)}$ (our choice in Section 3.1.2) is shown to be sensible compared to other possible values of c . What is plotted here is the empirical error rates of PCF-LDA classifier for varied values of scale parameter $c \in (0.1, 1000)$, shown as piecewise linear curves. The computed c is shown as vertical line segment. Different colors represent different realizations. Model (2) with $(p, N) = (200, 20)$ is used to generate set classification data.

C Comparison of different statistics (Section 4)

We report the results of our experiments in Table 3 leading that the use of $T(r)$ exhibited better performance than using other choices of statistics. For the comparison, we have considered the following statistics;

1. a modified sum of squared t -statistics $T_1^T(r) = \boldsymbol{\eta}_r'(D_r + \tau \mathbf{I}_m)^{-1} \boldsymbol{\eta}_r$, for $\tau \in [0, \text{trace}(S_r)]$ (Srivastava and Du, 2008),
2. a version of Hotelling's T^2 statistic $T_2^T(r) = \boldsymbol{\eta}_r'(S_r + \tau \mathbf{I}_m)^{-1} \boldsymbol{\eta}_r$,
3. Dempster's statistic $T_3(r) = \boldsymbol{\eta}_r' \boldsymbol{\eta}_r / \text{tr}(S_r)$ (Dempster, 1960, Srivastava and Fujikoshi, 2006), and
4. the ratio of between-group and within-group distances $R_1(r) = B_1(r)/W_1(r)$, where

$$B_1(r) = \sum_{k=1}^2 \sum_{i,j \in G_k} \rho_s(\mathcal{L}_i^{(r)}, \mathcal{L}_j^{(r)}), \quad W_1(r) = \sum_{i \in G_1} \sum_{j \in G_2} \rho_s(\mathcal{L}_i^{(r)}, \mathcal{L}_j^{(r)}).$$

Table 3 summarizes the performance of using these statistics in terms of the empirical error rates when r is selected by the corresponding statistic, and the classification method is chosen as either LDA, QDA or the 3-Nearest Neighbors. The empirical error rates when using $\hat{r}(T)$ are significantly smaller than those from using T_2 , T_3 and R_1 , and are comparable to those from using T_1 , for various dimension-sample size situations.

D Validation for the permutation test (Section 4)

We first demonstrate that the proposed permutation test leads to uniform null distribution for p-values. we report in Fig. 6 using an envelop QQ plot (Lee, 2007). The empirical power of the proposed permutation test is also examined. The power increases with the signal strength σ , and also with the sample size N , as shown in Figures 7 and 8. We used $p = 10$, $\sigma = 0, 1/2, 1, 3/2, 2$ with fixed $N = 20$ with the model (2) (cf. Section 5) for Figure 7, and $\sigma = 1$ with $N = 10, 20, 50, 100, 200$ for Figure 8. Note also that the power also increases with the number of instances n_i grows.

Table 3: The empirical error rates (mean and standard error based on 100 simulations) show that the performance of using T is superior to using others. The last row, labeled as ‘Base,’ is the minimum empirical errors given by exhaustive search over all $1 \leq r \leq R$. Model (2) in Section 5.2 is used with hyper-parameters $\delta = 1$, $\sigma = 3$, $\rho = 0$ and $\tau_s = t_s \text{tr}(S_r)$, $\{t_s\}_{s=1}^3 = \{0.01, 0.1, 0.5\}$.

(%)	$(p, N) = (20, 10)$			$(p, N) = (100, 10)$		
	LDA	QDA	3-NN	LDA	QDA	3-NN
T	6.92(0.54)	6.02(0.44)	9.32(0.63)	6.74(0.49)	6.65(0.45)	13.05(0.85)
$T_1^{T_1}$	5.76(0.48)	5.09(0.40)	8.04(0.54)	5.81(0.43)	5.82(0.41)	11.68(0.79)
$T_1^{T_2}$	4.52(0.37)	4.34(0.38)	6.88(0.46)	5.11(0.37)	5.20(0.37)	8.68(0.48)
$T_1^{T_3}$	6.28(0.44)	6.16(0.45)	9.62(0.55)	6.85(0.53)	6.94(0.53)	10.41(0.65)
$T_2^{T_1}$	12.23(0.50)	9.98(0.40)	14.93(0.64)	15.56(0.91)	14.35(0.79)	23.08(1.27)
T_3	12.23(0.50)	9.98(0.40)	14.93(0.64)	15.95(0.90)	14.70(0.79)	23.47(1.27)
R_1	12.23(0.50)	9.98(0.40)	14.93(0.64)	16.10(0.89)	14.85(0.78)	23.64(1.25)
Base	2.25(0.18)	2.06(0.18)	3.87(0.24)	2.45(0.16)	2.47(0.15)	5.56(0.29)

$(p, N) = (200, 20)$			
(%)	LDA	QDA	3-NN
T	3.20(0.23)	3.22(0.24)	10.45(0.67)
$T_1^{T_1}$	3.02(0.21)	3.02(0.22)	9.80(0.56)
$T_1^{T_2}$	2.53(1.94)	2.52(0.19)	8.15(0.35)
$T_1^{T_3}$	4.30(0.48)	4.37(0.49)	10.42(0.63)
$T_2^{T_1}$	8.04(0.64)	7.70(0.59)	17.85(1.20)
T_3	8.41(0.65)	8.05(0.59)	18.33(1.24)
R_1	8.44(0.65)	8.07(0.59)	18.33(1.24)
Base	1.67(0.14)	1.66(0.14)	6.85(0.33)




Figure 6: Quantile-quantile (QQ) envelop plot of the empirical p-values under the null hypothesis of the permutation test proposed in Section 4.2. The quantile-quantile (QQ) plots (shown as red curves) of the empirical p-values, in nine difference combinations of dimensions and sample sizes ($p = 10, 100, 400, N = 10, 20, 100$) are shown. We use the QQ envelope plot to understand the natural variation of the empirical quantiles (Lee, 2007). The green line in each panel of Figure 6 shows the theoretical quantiles from the uniform(0, 1) distribution. Hundred QQ plots from random samples of the same size are overlaid as blue curves. The red QQ curves are inside of the blue curve bundles, from which we conclude that the empirical p-values follow the uniform(0, 1) distribution.

When the true underlying model has no covariance difference between groups, the proposed permutation test greatly enhances the performance. Table 4 compares the classification results before (using $\hat{r} > 0$) and after updating \hat{r} by the decision of the test. The empirical error rates are computed using LDA, while using other classifiers exhibits similar behavior. The error rates are notably decreased after applying the permutation test. The target r is zero because we used a hierarchical model with no covariance difference. In this experiment, the \hat{r} is updated to zero in 92% out of 100 simulations (when T is used). Performances of using different statistics are also summarized in Table 4. While the benefit of using the test is clear, different choices of test statistic do not significantly affect the results.




Figure 7: Empirical power function of the permutation test (Section 4.2) on different signal strength $0 \leq \sigma \leq 2$. Model (2) with $(p, N) = (10, 20)$ is used.

E Definitions of modified linear discriminant analyses (Section 5)

The LDA and QDA used in the simulation studies and real data analysis are modified from the usual definition, for use in the high-dimensional, low-sample size situations. The LDA is notorious for its poor performance when $p \geq n$ (Bickel and Levina, 2004, Ahn and Marron, 2010). Both classification rules are dependent on the (pooled) sample covariance matrices. We simply replace the sample covariance matrix $\hat{\Sigma}$ with a shrinkage estimator $\hat{\Sigma} + \gamma \mathbf{I}$ for a small positive constant γ . We have used $\gamma = 0.01$ for all analyses in the paper. The idea of modified discriminant analysis dates back to Friedman (1989).

An anonymous referee has suggested to compare the performances of the proposed method with those of the modified linear discriminant analyses appeared in Srivastava and Kubokawa (2007) and Yata and Aoshima (2012). Here precise definitions of those methods




Figure 8: Empirical power function of the permutation test (Section 4.2) on different sample size $N = 10, 20, 50, 100, 200$ (the number of sets). Model (2) with $p = 10$, $\sigma = 1$ is used.

are given.

Consider binary classification with $Y = \pm 1$. With the pooled sample covariance matrix \mathbf{S} and the sample group means $\hat{\mu}_+$ and $\hat{\mu}_-$, Fisher's linear discriminant function is, for a new observation \mathbf{x} ,

$$\delta(\mathbf{x}) = \text{sign}\left(\mathbf{x} - \frac{\hat{\mu}_+ + \hat{\mu}_-}{2}\right)^T \mathbf{S}^{-1}(\hat{\mu}_+ - \hat{\mu}_-). \quad (13)$$

As discussed in Section 5.1, use of a modified linear discriminant function is a common

Table 4: The empirical error rates (mean and standard error based on 100 simulations) are significantly reduced when the permutation test is conducted. The true underlying model has no covariance difference (Model (1) in Section 5.2.

(%)	$(p, N) = (20, 10)$		$(p, N) = (100, 10)$		$(p, N) = (200, 20)$	
	before	after	before	after	before	after
T	18.68(1.12)	7.38(0.49)	24.48(1.20)	14.20(0.74)	21.75(1.00)	13.83(0.62)
$T_1^{T_1}$	21.13(1.09)	7.38(0.50)	26.14(1.17)	14.20(0.74)	25.79(0.95)	13.83(0.62)
$T_1^{T_2}$	26.02(0.84)	7.93(0.60)	34.89(0.84)	15.04(0.90)	35.03(0.61)	14.34(0.71)
$T_1^{T_3}$	29.16(0.66)	8.63(0.74)	40.45(0.59)	14.84(0.88)	36.80(0.59)	14.08(0.67)
$T_2^{T_1}$	18.14(1.25)	8.35(0.73)	22.70(1.22)	13.19(0.53)	17.58(0.86)	12.95(0.45)
T_3	18.19(1.25)	8.07(0.69)	22.85(1.23)	13.20(0.53)	17.47(0.86)	12.95(0.45)
R_1	18.19(1.25)	8.07(0.69)	22.90(1.22)	13.20(0.53)	17.47(0.86)	12.95(0.45)
Exhaustive	4.47(0.27)		11.45(0.49)		10.94(0.36)	
($r = 0$)	7.58(0.47)		13.74(0.42)		12.79(0.42)	

practice, by replacing \mathbf{S} by $\mathbf{S} + \gamma \mathbf{I}_p$, for an appropriate value of γ . We have used

$$\delta_\gamma(\mathbf{x}) = \text{sign}(\mathbf{x} - \frac{\hat{\mu}_+ + \hat{\mu}_-}{2})^T (\mathbf{S} + \gamma \mathbf{I}_p)^{-1} (\hat{\mu}_+ - \hat{\mu}_-),$$

with $\gamma = 0.01$.

Srivastava and Kubokawa (2007) have compared several linear discriminant methods and concluded that minimum distance empirical Bayes rule (MDEB) performs best among all methods considered under the high-dimensional, low-sample-size situation. The MDEB is obtained by choosing $\gamma = \text{trace}(\mathbf{S}) / \min(n, p)$, so that the classification function is

$$\delta_{MDEB}(\mathbf{x}) = \text{sign}(\mathbf{x} - \frac{\hat{\mu}_+ + \hat{\mu}_-}{2})^T (\mathbf{S} + \frac{\text{trace}(\mathbf{S})}{\min(n, p)} \mathbf{I}_p)^{-1} (\hat{\mu}_+ - \hat{\mu}_-).$$

Yata and Aoshima (2012) have considered a hard-thresholded covariance estimate \mathbf{S}_ω to replace \mathbf{S} in (13). Denote the eigen-decomposition of the sample covariance matrix as $\mathbf{S} = \sum_{j=1}^p \lambda_j \mathbf{e}_j \mathbf{e}_j^T$, where $\lambda_1 > \lambda_2 > \dots > \lambda_{n-1} > 0$ and $\lambda_j = 0$ for $j = n, \dots, p$. For

$$\omega = \min\left\{\frac{\text{trace}(\mathbf{S})}{p^{1/2}n^{1/4}}, \frac{\text{trace}(\mathbf{S})}{\min(n, p)}\right\},$$

the hard-thresholded covariance estimate of [Yata and Aoshima \(2012\)](#) is defined as

$$\mathbf{S}_\omega = \sum_{j=1}^{n-1} \max\{\lambda_j, \omega \left(\frac{\lambda_j}{\lambda_j - \frac{\sum_{i=j+1}^{n-1} \lambda_i}{n-j}} \right)\} \mathbf{e}_j \mathbf{e}_j^T + \sum_{j=n}^p \omega \mathbf{e}_j \mathbf{e}_j^T.$$

The corresponding classification function is

$$\delta_{YA}(\mathbf{x}) = \text{sign}(\mathbf{x} - \frac{\hat{\mu}_+ + \hat{\mu}_-}{2})^T \mathbf{S}_\omega^{-1}(\hat{\mu}_+ - \hat{\mu}_-).$$

The majority voting classifiers, denoted as 'Classifier'-MV, are constructed as a majority vote of the $\delta(\mathbf{x}_i)$, $i = 1, \dots, n$; for example,

$$\delta_{MDEB-MV}(\{\mathbf{x}_i\}) = \text{sign} \left(\sum_{i=1}^n \text{sign}(\mathbf{x} - \frac{\hat{\mu}_+ + \hat{\mu}_-}{2})^T \mathbf{S}_\omega^{-1}(\hat{\mu}_+ - \hat{\mu}_-) \right).$$

The weighted voting classifier is defined by

$$\delta_{MDEB-WV}(\{\mathbf{x}_i\}) = \text{sign} \left(\sum_{i=1}^n (\mathbf{x} - \frac{\hat{\mu}_+ + \hat{\mu}_-}{2})^T \mathbf{S}_\omega^{-1}(\hat{\mu}_+ - \hat{\mu}_-) \right).$$

F Additional tables for Section 5 (Numerical studies)

See [Tables 5](#) and [6](#).

Table 5: Empirical misclassification rates ($\rho = 0$).

Model (1) with $\rho = 0$						
(p, N)	(20,10)	(200, 10)	(400,10)	(20,20)	(200,20)	(400,20)
PCF-LDA	25.41(5.43)	33.13(3.49)	36.57(3.88)	20.15(5.01)	28.34(3.82)	33.18(3.86)
PCF-QDA	29.07(6.00)	34.11(3.99)	37.06(3.80)	32.62(6.05)	31.84(3.94)	34.84(3.85)
PCF-SVM	22.03(5.45)	33.10(3.51)	36.61(3.88)	20.74(5.08)	28.30(3.81)	33.22(3.89)
PCF-DWD	18.75(4.28)	32.74(3.73)	36.60(4.01)	15.29(3.68)	26.74(3.34)	32.61(3.97)
PCF-MDEB	18.26(4.42)	32.59(3.82)	36.62(4.02)	14.81(3.65)	26.65(3.41)	32.59(3.93)
PCF-YA	21.32(7.64)	46.87(5.52)	47.65(3.84)	14.23(3.61)	41.38(5.97)	45.88(4.34)
LDA-MV	20.21(4.64)	38.46(5.01)	44.53(3.89)	16.63(4.14)	31.43(4.64)	38.69(4.10)
LDA-WV	18.18(4.60)	37.57(4.91)	44.23(3.80)	14.29(3.89)	29.44(4.92)	37.61(4.47)
QDA-MV	27.93(5.84)	49.63(1.32)	49.26(2.00)	21.40(5.30)	47.39(2.96)	49.48(1.53)
QDA-WV	26.31(5.77)	49.43(1.75)	49.41(1.56)	19.38(5.35)	47.12(2.81)	49.57(1.46)
SVM-MV	20.72(4.56)	38.27(4.36)	40.70(4.13)	16.82(4.29)	32.39(4.16)	38.27(4.70)
DWD-MV	19.55(4.14)	34.06(3.50)	37.58(3.70)	16.12(3.84)	28.30(3.73)	34.09(3.84)
MDEB-MV	21.11(5.95)	35.18(3.86)	38.86(4.08)	17.10(4.29)	28.99(4.11)	35.20(4.23)
MDE-WV	19.57(6.08)	33.69(4.37)	37.93(4.29)	14.71(4.07)	27.63(4.20)	33.42(4.01)
Model (2) with $\rho = 0$						
PCF-LDA	4.06(3.58)	2.35(3.84)	5.38(4.00)	0.85(0.85)	0.79(0.71)	2.22(1.25)
PCF-QDA	7.32(5.79)	2.30(3.77)	5.34(4.00)	3.82(3.77)	0.72(0.69)	2.16(1.24)
PCF-SVM	5.20(4.39)	2.35(3.84)	5.38(4.00)	1.57(2.00)	0.79(0.71)	2.22(1.25)
PCF-DWD	6.56(4.76)	2.35(3.93)	5.38(4.00)	3.01(2.85)	0.77(0.68)	2.22(1.24)
PCF-MDEB	6.14(4.60)	2.36(3.87)	5.38(4.01)	2.20(2.36)	0.77(0.69)	2.22(1.26)
PCF-YA	9.12(7.52)	22.47(10.64)	30.61(8.84)	1.50(1.59)	19.39(10.66)	26.43(9.89)
LDA-MV	39.78(4.26)	47.69(3.53)	46.24(3.40)	36.41(4.22)	45.91(3.49)	47.60(3.36)
LDA-WV	38.18(4.72)	47.48(3.49)	45.92(3.44)	34.79(4.42)	45.27(3.94)	47.80(3.66)
QDA-MV	4.71(2.14)	25.73(3.78)	24.09(3.17)	2.25(1.24)	35.07(4.25)	29.86(3.61)
QDA-WV	2.04(1.49)	21.96(3.84)	20.10(2.99)	0.63(0.64)	33.36(4.93)	27.40(3.71)
SVM-MV	39.95(4.80)	46.08(3.45)	44.99(3.54)	36.38(4.03)	46.48(3.86)	46.55(3.37)
DWD-MV	36.74(4.61)	38.75(4.19)	40.39(4.07)	34.77(4.08)	36.28(3.80)	38.93(3.41)
MDEB-MV	36.16(6.28)	42.81(3.64)	43.30(4.03)	32.23(5.47)	41.09(3.88)	43.46(3.89)
MDE-WV	34.34(7.26)	41.49(3.92)	42.77(3.99)	30.29(6.07)	39.91(4.23)	43.31(3.31)
Model (3) with $\rho = 0$						
PCF-LDA	5.80(4.07)	17.90(6.87)	27.88(7.40)	2.84(1.89)	12.05(5.52)	21.03(6.94)
PCF-QDA	6.25(4.53)	18.37(7.12)	28.19(7.27)	4.14(3.15)	13.99(6.58)	22.71(7.60)
PCF-SVM	5.75(3.83)	17.89(6.86)	27.40(5.78)	3.17(2.18)	11.97(5.43)	21.03(6.94)
PCF-DWD	6.10(3.79)	18.10(6.58)	28.14(4.86)	3.59(2.55)	11.97(5.15)	20.88(6.74)
PCF-MDEB	7.34(5.69)	20.00(6.90)	30.10(6.38)	3.49(2.46)	12.50(4.82)	23.35(7.18)
PCF-YA	11.44(8.05)	41.13(8.31)	44.94(7.57)	3.61(2.76)	39.85(6.84)	44.71(6.09)
LDA-MV	36.45(6.03)	36.31(4.42)	36.27(4.29)	35.04(5.54)	35.61(4.58)	34.94(3.58)
LDA-WV	36.10(6.33)	35.92(4.36)	36.25(4.27)	34.38(5.85)	35.52(4.53)	34.67(3.48)
QDA-MV	27.64(4.25)	32.46(3.33)	34.20(3.96)	28.66(6.94)	33.78(3.65)	34.89(3.34)
QDA-WV	25.59(4.27)	31.28(3.30)	33.24(3.88)	26.29(6.85)	32.28(3.39)	34.01(3.59)
SVM-MV	36.74(6.14)	33.48(4.45)	34.74(4.41)	33.13(6.32)	32.75(4.93)	32.57(3.66)
DWD-MV	29.35(6.50)	28.72(4.04)	31.80(4.26)	28.99(6.07)	24.65(3.80)	27.87(3.57)
MDEB-MV	35.61(11.43)	31.01(6.47)	33.96(5.81)	37.45(12.18)	26.94(6.37)	30.39(6.59)
MDE-WV	34.81(12.20)	30.52(6.48)	33.74(5.86)	36.89(12.75)	26.30(6.50)	30.14(6.78)
Model (4) with $\rho = 0$						
PCF-LDA	13.78(5.66)	33.40(4.06)	38.07(4.07)	1.53(1.56)	30.31(4.03)	32.49(3.58)
PCF-QDA	21.14(5.68)	33.50(3.81)	38.38(4.15)	17.10(4.75)	33.20(4.25)	34.75(3.80)
PCF-SVM	15.53(5.41)	33.42(4.03)	38.09(4.05)	4.95(4.14)	30.27(4.02)	32.50(3.55)
PCF-DWD	15.73(4.81)	32.75(4.28)	37.82(4.15)	10.21(3.54)	29.02(3.78)	31.83(3.56)
PCF-MDEB	15.36(4.55)	32.65(4.19)	37.80(4.05)	8.80(3.77)	28.91(3.77)	31.65(3.53)
PCF-YA	16.57(8.63)	45.65(6.03)	48.23(4.64)	0.54(0.97)	42.89(5.52)	46.91(4.51)
LDA-MV	38.90(4.02)	41.60(3.85)	45.82(3.87)	35.52(4.39)	38.02(4.10)	39.27(4.41)
LDA-WV	37.25(4.36)	41.12(4.21)	45.45(4.11)	33.84(4.70)	37.50(4.11)	38.24(4.94)
QDA-MV	4.34(1.99)	49.68(1.46)	49.23(1.99)	2.72(1.54)	47.28(2.96)	49.61(1.84)
QDA-WV	1.02(1.25)	49.57(1.57)	49.53(1.66)	0.23(0.35)	47.02(3.30)	49.63(1.40)
SVM-MV	38.28(4.20)	41.38(3.98)	41.50(3.89)	35.59(4.15)	39.51(4.36)	38.72(4.39)

Table 6: Empirical misclassification rates ($\rho = 0.5$).

(p, N)	Model (1) with $\rho = 0.5$					
	(20,10)	(200, 10)	(400,10)	(20,20)	(200,20)	(400,20)
PCF-LDA	23.69(5.49)	33.76(4.58)	38.77(4.38)	18.60(4.52)	28.53(4.68)	33.93(4.75)
PCF-QDA	27.15(5.66)	36.42(4.25)	40.14(4.31)	30.98(6.45)	33.52(4.44)	37.46(4.30)
PCF-SVM	20.57(4.96)	33.73(4.58)	38.75(4.38)	19.03(4.99)	28.34(4.68)	34.01(4.81)
PCF-DWD	17.44(3.96)	34.31(4.72)	39.58(4.42)	14.27(3.46)	27.78(4.98)	34.28(4.76)
PCF-MDEB	17.09(3.99)	34.34(4.82)	40.04(4.58)	13.94(3.31)	27.94(4.96)	34.41(4.82)
PCF-YA	19.35(6.27)	45.50(5.28)	48.62(4.02)	16.16(4.89)	42.35(5.91)	46.80(4.08)
LDA-MV	18.08(3.99)	36.82(5.32)	43.51(3.91)	14.53(3.86)	29.71(5.01)	37.13(4.41)
LDA-WV	16.56(4.07)	36.17(5.08)	43.17(4.06)	12.99(3.70)	28.43(5.02)	36.13(4.51)
QDA-MV	24.11(4.96)	49.19(1.94)	48.87(2.49)	18.54(4.97)	45.96(3.24)	49.24(2.02)
QDA-WV	22.75(5.01)	48.94(2.18)	48.98(2.14)	16.55(4.87)	45.82(3.55)	49.18(2.11)
SVM-MV	18.39(3.76)	35.89(4.53)	38.74(4.29)	14.53(3.81)	32.04(4.67)	36.08(4.12)
DWD-MV	17.17(3.74)	32.33(4.07)	37.69(3.60)	13.78(3.30)	26.93(3.61)	34.47(4.11)
MDEB-MV	18.77(5.01)	32.73(4.20)	36.84(4.01)	14.71(3.79)	26.41(4.04)	32.58(4.10)
MDE-WV	17.31(5.24)	31.33(4.05)	36.23(4.05)	12.95(3.74)	25.39(3.92)	31.39(3.96)
(p, N)	Model (2) with $\rho = 0.5$					
	(20,10)	(200, 10)	(400,10)	(20,20)	(200,20)	(400,20)
PCF-LDA	3.91(4.23)	1.66(2.70)	3.66(2.87)	0.49(0.63)	0.51(0.62)	1.58(1.05)
PCF-QDA	6.77(5.22)	1.69(2.78)	3.67(2.84)	2.64(2.46)	0.49(0.61)	1.55(1.09)
PCF-SVM	4.68(4.55)	1.66(2.70)	3.66(2.87)	0.88(1.03)	0.51(0.62)	1.58(1.05)
PCF-DWD	6.16(4.95)	1.67(2.74)	3.62(2.79)	2.09(1.65)	0.50(0.61)	1.56(1.05)
PCF-MDEB	5.71(4.81)	1.67(2.74)	3.64(2.80)	1.35(1.28)	0.51(0.61)	1.58(1.05)
PCF-YA	10.15(8.05)	19.26(10.01)	26.92(9.30)	1.28(2.02)	18.82(11.20)	23.74(9.72)
LDA-MV	40.53(4.46)	47.84(3.49)	46.83(3.64)	36.98(4.04)	46.53(3.66)	48.24(3.69)
LDA-WV	39.52(4.66)	47.54(3.30)	46.68(3.64)	35.87(4.31)	46.18(3.69)	48.24(3.46)
QDA-MV	3.44(1.82)	22.50(3.79)	21.30(3.01)	1.59(0.94)	31.51(4.88)	26.50(3.95)
QDA-WV	1.30(1.21)	18.12(3.61)	17.01(2.97)	0.34(0.38)	28.94(5.69)	23.67(3.75)
SVM-MV	40.25(4.94)	46.33(3.49)	45.69(3.62)	37.31(3.77)	47.06(3.84)	47.00(3.31)
DWD-MV	37.27(4.63)	38.63(4.15)	39.86(4.13)	35.06(3.73)	36.31(3.58)	38.90(3.37)
MDEB-MV	36.18(6.17)	42.69(3.69)	43.23(3.98)	32.50(5.18)	41.42(4.01)	43.39(3.49)
MDE-WV	34.06(7.11)	41.88(3.92)	43.21(3.71)	29.82(5.89)	40.64(4.26)	43.45(3.35)
(p, N)	Model (3) with $\rho = 0.5$					
	(20,10)	(200, 10)	(400,10)	(20,20)	(200,20)	(400,20)
PCF-LDA	5.55(4.46)	17.37(6.91)	27.01(7.95)	2.50(1.94)	11.60(5.58)	20.40(7.09)
PCF-QDA	6.15(4.99)	18.01(7.26)	27.33(7.93)	3.64(3.10)	13.87(6.63)	22.25(7.74)
PCF-SVM	5.46(4.14)	17.35(6.90)	26.57(6.52)	2.79(2.04)	11.57(5.45)	20.46(7.13)
PCF-DWD	5.96(4.26)	17.35(6.71)	27.17(5.45)	3.36(2.48)	11.36(5.09)	20.19(6.96)
PCF-MDEB	7.38(7.21)	18.67(6.47)	29.07(6.34)	3.18(2.31)	11.62(4.87)	22.38(6.07)
PCF-YA	11.48(9.11)	40.59(8.34)	44.36(7.29)	3.27(3.00)	39.45(6.67)	45.09(5.57)
LDA-MV	37.96(6.15)	36.93(4.64)	36.82(4.15)	36.06(5.55)	36.54(4.74)	35.38(3.49)
LDA-WV	37.28(6.41)	36.69(4.49)	36.58(4.10)	35.27(6.13)	36.29(4.67)	35.19(3.54)
QDA-MV	26.55(4.12)	32.31(3.32)	33.99(3.89)	27.33(6.55)	33.33(3.67)	34.68(3.40)
QDA-WV	24.04(4.21)	30.73(3.26)	33.03(3.89)	24.25(6.29)	31.58(3.45)	33.68(3.66)
SVM-MV	37.96(6.37)	34.06(4.57)	34.87(4.59)	34.42(6.13)	33.76(4.60)	32.98(3.81)
DWD-MV	30.16(6.76)	28.80(4.00)	31.65(4.37)	29.26(6.13)	24.68(4.20)	27.99(3.67)
MDEB-MV	35.40(11.35)	31.05(6.46)	34.06(5.79)	37.01(12.15)	27.18(6.03)	30.54(6.56)
MDE-WV	34.66(11.99)	30.52(6.47)	33.78(5.83)	36.53(12.83)	26.50(6.16)	30.16(6.70)
(p, N)	Model (4) with $\rho = 0.5$					
	(20,10)	(200, 10)	(400,10)	(20,20)	(200,20)	(400,20)
PCF-LDA	1.97(1.70)	34.18(4.57)	39.33(4.51)	0.15(0.29)	29.73(3.58)	33.32(4.42)
PCF-QDA	6.81(2.86)	36.02(4.76)	41.06(3.99)	2.98(1.54)	33.95(4.11)	37.14(4.02)
PCF-SVM	2.86(2.10)	34.21(4.58)	39.31(4.51)	0.48(0.75)	29.71(3.55)	33.34(4.40)
PCF-DWD	4.46(2.77)	34.34(5.06)	39.89(4.58)	1.85(1.72)	29.56(3.85)	33.89(4.88)
PCF-MDEB	3.76(2.29)	34.59(5.40)	40.09(4.75)	0.99(1.04)	29.51(3.82)	33.95(5.01)
PCF-YA	5.08(5.85)	45.86(5.66)	47.98(4.66)	0.67(1.77)	43.20(6.01)	46.29(4.63)
LDA-MV	39.29(3.86)	41.80(3.51)	44.54(4.08)	36.17(4.39)	38.78(3.83)	38.55(4.62)
LDA-WV	38.01(3.96)	41.12(4.05)	44.44(4.17)	34.84(4.42)	37.92(4.15)	37.83(4.73)
QDA-MV	3.25(1.68)	49.17(1.90)	49.28(2.31)	1.89(1.18)	45.65(3.12)	49.49(1.76)
QDA-WV	0.55(0.90)	49.07(2.10)	49.22(1.96)	0.10(0.23)	45.32(3.70)	49.37(1.61)
SVM-MV	38.46(4.31)	40.39(3.67)	40.08(4.36)	35.50(4.34)	40.23(4.16)	37.48(4.40)